STScI-NGST-R-0001A

Space Telescope Science Institute
Next Generation Space Telescope Mission

Operations Concept Interim Report: Issues and Considerations for NGST Science Operations

24 January 2000

Issue A

REVISION HISTORY

ISSUE	DESCRIPTION	DATE
v7	Initial release -- First official version after 13 December 1999 review	4-Jan-00
v8	Draft for review prior to sign-off	7-Jan-00
A	Release version	24-Jan-00

Next Generation Space Telescope Mission

Operations Concept Interim Report: Issues and Considerations for NGST Science Operations

January 24, 2000

PREPARED BY: Henry C. Ferguson________________ ____________
NAME ORG.

____________________________________ ____________
SIGNATURE DATE

____________________________________ ____________
NAME ORG.

____________________________________ ____________
SIGNATURE DATE

APPROVED BY: ____________________________________
SIGNATURE

____
NAME

____
TITLE

NGST Operations Concept Interim Report: Issues and Considerations for NGST Science Operations

Edited by H. C. Ferguson

The NGST Operations Concept Working GroupMembers

Stefano Casertano
Space Telescope Science Institute
Rodger Doxsey
Space Telescope Science Institute
Harry Ferguson (editor)
Space Telescope Science Institute
Jonathan Gal-Edd
Goddard Space Flight Center
Bob Hanisch
Space Telescope Science Institute
John Isaacs
Space Telescope Science Institute
Wayne Kinzel
Space Telescope Science Institute
Melissa McGrath
Space Telescope Science Institute
Alexander Storrs
Space Telescope Science Institute
Meg Urry
Space Telescope Science Institute

Contents

NGST Operations Concept

1 Introduction

The Next Generation Space Telescope is envisioned as an international facility, with competitively selected observing programs spanning a wide range of topics. The Space Telescope Science Institute will serve as the science operations center, with responsibilities that include overseeing the proposal selection process, scheduling the observations, characterizing and calibrating the instruments, and providing the archive of reduced and calibrated data, as well as planning and conducting operations of the spacecraft.

The cost to operate NGST represents a significant fraction of the overall life cycle budget. Early planning for science operations can drive down the cost of the mission by allowing considerations of operations complexity and efficiency to have a bearing on the NGST design. As the design of the observatory and planning for the mission progress, the operations concepts will also undoubtedly evolve, and feed back into further iterations on the design.

In March 1999, a diverse group of staff members from STScI and GSFC was assembled to begin assessing various operations concepts and tradeoffs. Starting with the broadest definition of the NGST architecture, a passively cooled 8m NIR/MIR telescope at L2, the group tried to identify and characterize the major factors that could affect NGST operations such as viewing geometry, sky accessibility, background environments, etc. The study also explored various concepts for the science mix of observations (GO, GTO, key projects, and legacy projects) and different concepts for scheduling (block vs. queue scheduling).

This interim report of the Operations Concept study presents a high-level overview of the major issues that affect NGST science operations. Our primary goal is to describe an operations philosophy, identify key issues for trade studies, and elicit feedback, rather than propose a single set of solutions.

1.1 Scope of the Operations Concept Study

We define NGST operations to include all aspects of the observatory that affect, guide, enable or limit the ability of NGST to acquire data and the scientist's or engineer's access to that data. This scope includes all flight-system and ground-system software, the techniques and frequency of flight/ground interaction, and the balance between autonomous operations and skilled human planning and engineering analysis. The study includes an examination of NGST observing constraints and overheads, and the implications of different observing strategies and science scheduling policies on operations.

The Operations Concept study is one of several studies relevant to operations, and its exact charter has been evolving as the study progresses. As mentioned above, the current goal has been to identify the key operations issues that have a bearing on NGST design, lifetime costs, and scientific productivity. The study will evolve into a more detailed specification of an operations plan, with a more detailed specification of ground/flight interactions, observing modes, different kinds of spacecraft activities, different mission phases, and observation scheduling. A more detailed report containing this kind of information (already addressed at different levels of detail in some of the previous NGST operations studies), as well as results of further study of the issues presented here, is due in mid-2000. The intent is to provide a definitive (but still relatively high-level) baseline description of how NGST will be operated.

The Operations concept study has been informed by a number of other studies, some of which have gone on in parallel. We have attempted to incorporate some of the intermediate results of these studies in our report where relevant. For example, the discussion of the event-driven on-board scheduling and the architecture that is envisioned to support it, reflects the work of the Software Operations Working Group, and is included in this document to provide a snapshot of the design concepts currently being considered and how they bear on operations.

1.2 Overview of Findings and Recommendations

Based on current understanding of the NGST design (mostly the Yardstick concept) and the various constraints and limitations on NGST observing, the following appear to be the issues that could have the most important bearing on science productivity and operational complexity:

1. Data Volume
2. OTA stability and calibration
3. Thermal crosstalk in the ISIM
4. Overheads for small slews & dithering
5. Telescope boresight roll restrictions
6. Guide camera
7. Thermal stability and time to settle after slews
8. Momentum management
9. Radiation environment

It is too early to make specific recommendations for many of these issues. Not all of the technical issues have been studied by the operations group, and for at least some the Yardstick design may no longer be the most useful benchmark to consider. There are, however, a few clear recommendations that can be proposed at this time.

Recommendations:

1. For at least the first year of science operations, NGST should have the capacity to downlink data without on-board cosmic-ray rejection. We estimate this requires a downlink capacity of 250 Gbits/day with sufficient on-board storage to hold the data between downlinks.
2. The overheads for slews and thermal and vibrational settling should be kept small enough that they do not drive operations. Typical NGST observing will require of order 1 large-angle maneuver per day and 80 small-angle maneuvers (on a scale of 10-20 arcsec) for dithering. The total overhead for these motions should be less than ~2 hours per day.
3. The combined total of overheads, including those for momentum management, wavefront calibration, detector readout, and instrument reconfigurations and slews, must be kept low (10-20%), and any individual component should be kept small enough that it does not have to be a key parameter in the scheduling algorithm.
4. The ISIM should be designed to allow parallel operations of instruments without thermal or electronic crosstalk between them.
5. Increasing NGST boresight roll flexibility could greatly simplify scheduling. If the off-nominal roll capability for most targets is less than ~10 degrees on any given calendar day, the choice of field orientation for a specific observation will become an operations choice rather than a scientific decision.
6. The use of the camera as a guider, combined with restrictions in orientation, will occasionally conflict with the scientific desire to observe fields with multiple instruments at specific orientations. Further study of the practical implications of using the camera as a guider is needed.
7. The design and complexity of the scheduling system will depend in part on the philosophy for rescheduling observations that are missed or degraded due to particle radiation or other problems. Early policy decisions in this area are needed.

1.3 Differences Between NGST and HST Affecting Operations

While both NGST and HST are international facilities supporting a large number of observers and a large suite of observations, there are some key differences between the two observatories that will lead to a different (and to a large extent much simpler) operations strategy for NGST. Table 1 gives an overview of the major operational differences between NGST and HST.

HST	*NGST*	*Implications*
Low earth orbit	L2 halo orbit	NGST observations not interrupted by earth occultation. Higher observing efficiency can be attained, but overhead activities such as slews and calibrations (often scheduled during occultation for HST) will have a more explicit impact. Accurate orbit position not required for NGST
Servicing possible	Servicing not possible	Lifetime issues will be a bigger concern for NGST. Operations with partially-working instruments/control systems will have to be considered.
Data relay through TDRS	Data relay through dedicated ground station(s)	NGST will have less competition for downlink resources allowing more consistent scheduling of data downlinks Long contacts permit downlink of large data volume, but long periods without contact require more spacecraft autonomy and impose scheduling requirements for real-time operations.
OTA very stable	OTA stability a concern	Thermal stability should be provided by active thermal control; otherwise complex scheduling to restrict pitch angle changes between observations will be required to avoid large thermal settling time. Periodic wave-front measurement will be required, and regular adjustments might be necessary after orbit maintenance maneuvers.
SAA passages ~7 orbits per day	No SAA	More scheduling flexibility
Magnetosphere shields HST from most particle storms.	Solar flares/coronal mass ejections are a major source of particle background	NGST observation plan execution needs to be interruptible to respond to increases in particle background. On-board monitoring should be used to suspend observation plan execution when solar radiation is high, and possibly terminate observation plan execution if dangerous levels are reached.
Target windows determined largely by solar & lunar avoidance	Target windows determined largely by sunshade constraints	NGST targets typically available for two blocks of time, with duration and available roll determined by ecliptic longitude & latitude

HST	***NGST***	***Implications***
Solar array geometry typically allows ±30 deg of boresite roll flexibility	Sunshield geometry restricts boresite roll to a small range of angles	Observing programs requiring many days at fixed roll may be infeasible. Scientific compromises will be necessary to allow all programs to schedule.
Limited onboard computing capability programmed in assembly language	Modern flight processors and programming languages	Much more flexibility and autonomy can be built into the flight software. Offers potential for on-board data processing.
Short slew time	Long slew time	NGST may have long slew times due to limited capacity of reaction wheels, with a significant impact on observing efficiency. High capacity reaction wheels with isolation mounts are desired, which provide fast slews but also provide jitter-free maneuvers of one or two FOV.
Momentum dumps via magnetic torquers	Momentum dumps via thrusters	Conserving propellant may be an issue in NGST scheduling. Long observations may be interrupted for momentum dumps due to limited capacity of reaction wheels. Optimal sunshield geometry will provide torque neutral attitude; scheduling attitudes that alternate about the torque neutral attitude will conserve propellant.
Separate guiders with a large field of view, and more accurate spatial response than camera.	Guider may be part of the camera.	Location of guide stars may influence observing/dithering strategies. With guide stars close to target, roll error (about 1 arc-second) will have minimal effect on image stability.
Warm	Cold	The requirement to maintain a cold, stable temperature may impose restrictions on instrument operations and observing strategies. But the telescope needs to be cold to observe in the infrared.
Orbit relatively stable, with slow decay over several years.	Orbit not stable.	Orbit maintenance for HST required after several years, done by Space Shuttle during servicing. Orbit maintenance for NGST required frequently (perhaps monthly), done by thrusters and may require OTA calibration afterwards.
Small-format detectors	Large-format IR detectors	Data volume may be significantly higher for NGST than for HST. Most observations will not require target acquisitions to support small apertures.

1.4 Structure of this Document

This interim report is broken into seven sections. The first provides a general overview. The second section establishes a general operations philosophy, and tries to lay out a basic operations framework that will support a reasonable universe of options that could arise as decisions are made about the NGST architecture and science program. The third section describes the core "nugget" of flight and ground systems functionality. It broadly defines an architecture that easily allows either quasi-real-time or completely preplanned operations, and outlines the elements that will allow for incorporation of on-board interruptions and decisions. The fourth section describes various technical aspects of NGST operation that are yet to be determined and will result in adjustments of the operations concept within the full range of possible concepts. The topics include OTA stability, OTA alignment requirements, momentum management, solar flares, slew settle times, ground stations and contact frequency, and instrument warmup/stability questions. A final portion of this section explores the relationship between timescales for various spacecraft parameters (e.g. OTA stability) and the science program (in terms of number of targets/year). The fifth section describes various aspects of the science program. The purpose (in this interim report) is to cover a range of possibilities, discuss the implications of them, and prompt further discussion within the STScI, the NGST community, and the broader astronomical community. As a part of this discussion, we try to identify factors that may have particularly large cost, as well as scientific impact. The data pipeline and archive is discussed in section 6, and the summary and goals for future work are discussed in the final section.

2 Approach to the Development of an NGST Operations Concept

At this early phase of the project, the Operations Concept cannot be defined in great detail, any more than the hardware design can be described in great detail. Defining and adopting a detailed Operation Concept at this stage could place unacceptable burdens on the costs of hardware development. This would be inappropriate for a mission that will have many great technical challenges. Rather, the Operations Concept should start with a basic framework that can accommodate the range of technical and scientific decisions yet to be made.

Fundamentally, operating NGST should be viewed as the problem of using a ground computer system to operate a flight computer system. While the mirror and optics are the heart of the system to the opticians and astronomers, from an operational perspective they are not much more than a peripheral device. The fundamental nature of operations will consist of sending instructions up to the completely computerized NGST, and interpreting the data provided by the flight software. The systems on the ground which carry out these functions will need to be highly automated, both to deal with the inherent complexity of the task and to reduce the personnel costs of the operations phase. The ground and flight systems must also have flexibility so that operations staff can use them effectively for a wide range of tasks, including dealing with unanticipated problems. It is natural, therefore, that the initial Operations Concept should focus on developing an integrated philosophy for the flight software, the ground software, and the interface between them. This core concept must be sufficiently robust so that it can accommodate the range of future decisions concerning both the hardware and the science program and still maintain its validity. Overall goals in this process are to keep lifecycle costs low, maintain a high level of automation where appropriate, incorporate common hardware and software architecture at all levels in the system, using standard components as much as possible, and keep the overall system simple and operator-friendly.

As the design of the hardware, including the OTA, the instruments, and the spacecraft support components progress, the core operations concept will be used to help evaluate and select among trade-off options. In some cases, the life-cycle operations costs will be a significant factor in the trade-off and may dominate the decision process. In other cases the life-cycle costs may be neutral or minor in the decision making process. As these decisions are reached the Operations Concept will naturally expand to cover a larger scope with more specific details. These decisions will then drive further design decisions for both the ground and flight software systems.

Similarly, the NGST community and the astronomy community as a whole will be making fundamental decisions about the nature of the NGST science program. The Operations

Concept will be used to help evaluate and select among options. In some aspects, the life-cycle operations costs will be a significant factor in a decision. In other aspects the operations costs may be a neutral or minor factor. As with the hardware decisions, once science program decisions are made the Operations Concept and the requirements on the ground and flight systems will become more specific, narrowing the range of options previously held open.

Early definition of the core flight/ground systems Operations Concept will allow for an early start on the most fundamental aspects of both systems. The basic operational interactions and interfaces between the ground and flight systems can and should be defined long before the details of each specific operation and piece of hardware to be controlled are known. This will allow for the early setting of requirements on how each element of the flight system will be operated from the ground. Further, this process should clearly identify the operational requirements the hardware systems must meet in order for the Operations Concept to remain viable. The early definition of these concepts will allow for early definition of test cases that reflect the intended operation of each element. The early development of the fundamental commanding and telemetry aspects of the integrated flight/ground system design will also allow for the early development of Integration-and-Testing versions of the system, to be used in an operational manner while testing hardware elements. The operations concepts will need to adapt as decisions are made concerning the instruments, the telescope, the spacecraft, and other elements of the overall system, but also operations cost and complexity must be included among the design tradeoffs. Some of these decisions will be taken before there is any substantial investment in the flight and ground operations systems, but most will be made later in the program. Some of the late decisions that impact operations will be the result of “descopes” made during the hardware development phase, some will result from the discovery that the technology is not as mature as planned, and some will result from measurement of actual on-orbit performance. A common flight and ground operations strategy that is flexible enough to cover a wide range of potential operations scenarios will most effectively be able to respond to the specific problems which develop later in the program

3 Core Concept for NGST Flight and Ground Systems

The following attributes are considered to be necessary for the suite of ground and flight systems that must support NGST operations.

- The design must be flexible enough to allow real-time, quasi-real-time, and completely preplanned command and control of spacecraft activities.
- There should be a consistent design philosophy among the different systems that must interact to operate the observatory.
- The on-board activities should generally be event driven rather than tied to a specific time, although the capability to tie activities to specific times must exist. The syntax should be compatible and the ground system should be capable of producing either an event-driven timeline or a timeline with times of the activities explicitly specified.
- The systems and their interactions must be simple and flexible enough to allow operations strategies to be conveniently revised based on in-flight experience.

The remainder of this section deals mostly with the event-driven control system envisioned to meet these broad goals. Much of the discussion deals with the on-board command and control system. This is where most of the advanced effort is needed, both because the hardware must be acquired much earlier than the ground-system hardware, and because careful planning of the on-board system can result in major simplifications to the ground system and to the operations activities.

3.1 Flight Software Role in Operations

The flight software shall be designed to carry out all normal activities of the spacecraft, instruments, and support equipment, either in a continuous fashion or as the result of a specific request from ground controllers. The flight software will directly manage all on-board hardware systems, taking direction from the uplinked observing sequence for major activities such as slewing to a new target, acquiring guide stars, and executing a sequence of observations with an instrument. Routine operations will require uplink of information and parameters needed for major tasks, not detailed or repetitive information needed for commanding each individual mechanism. During routine observing, the flight software will manage the low level commands that go directly to the hardware. The operations interface will not need to include low-level hardware commanding. The flight software will also handle the appropriate interfaces between on-board systems. It will not require uplinked commands to configure two systems that routinely communicate; this should be handled on-board by the software managing the systems. From the perspective of operations, the flight system should have the “look and feel” of a single on-board system, no matter how many different computers, operating systems, and languages were used. The flight software will therefore have a consistent and robust design philosophy for the packets of data that are sent to it to initiate actions. This commonality of design of command packets will cross all on-board hardware and software systems, in human-readable formats to improve testability pre-launch and visibility into commanded operations post-launch. The data sent to the flight software will generally be in a high level form in engineering units. The on-board system will not require operations to send the same data (e.g. target position) to multiple processes on-board (e.g. slew calculations, sun avoidance calculations, differential aberration calculations), but rather will only require uplink of the minimal amount of data and will distribute it on-board.

The flight software system will provide reports of system actions that are downloaded to the ground upon initiation of communications contact. These reports will provide an indicator of the observation plan execution status, an observation log of observations executed or skipped, and an event log of requested and completed activities including start and stop times and completion status (e.g. success, failure, time-out).

3.2 Stored and/or Real-Time Control

Substantial commonality can be achieved by using the same basic format for uplink data for both direct, real-time commanding and for commands stored onboard for execution at a later time. The format can include flags or markers that allow this distinction, but fundamentally the information provided can be identical. This will also allow for substantial commonality in the ground system elements and the I&T systems which must construct and validate the command packets. It will also provide for easier evolution of routine functions from test systems to real-time operational use and then to regular stored command usage.

A single design will support both a completely real-time commanding operation and a completely stored command operation. This flexibility is needed to ensure support of a wide range of operational scenarios, both pre-launch and post-launch. Much of the initial testing in the pre-launch era will be via real-time test scenarios. It should be easy to then package up these scenarios and repeat them from stored commanding in later tests, saving time and providing for easy regression testing. We can expect that some post-launch operational scenarios will start out as real-time scenarios and then be transitioned into stored commanding when they have matured. As an example, scenarios used for aligning elements of the OTA are likely to be in this category. This capability in the flight/ground interface will allow tailoring the actual operations scenarios to suit the situation we find ourselves in after NGST is deployed, checked out, and operational. If the OTA and SIs are very stable, then moving to a completely stored command scenario may reduce the cost of operations. On the other hand, if the OTA requires frequent and unpredictable intervention to maintain alignment, then a real-time component of support may be required long after launch.

3.3 Sequential, Relative Time Operations

The flight software system will need an executive process that allows the execution of the commanded activities in a sequential manner, rather than requiring absolute times on the commands. The L2 orbit of NGST will provide an observing site without periodic natural interruptions, such as occultations on the time scale of a typical observation. Efficiency can be maximized if the system allows for the next task to start as soon as the prior task has finished. Sequential execution will allow for tasks of variable duration without requiring the ground system to predict exact completion times, or allocate worst-case execution times. The corollary to this is that the software and hardware must be designed so that tasks signal completion (and error codes) to the executive process. The flight software system may need to support fixed time execution of a timeline to support launch and deploy operations, as well as possible astronomical observations with critical time constraints. In the case of interspersed event-driven and time-critical observations, some inefficiency can be tolerated to ensure that the time-critical observations execute as scheduled.

3.4 Pausing and Jumping

The flight software system will provide a capability to pause and resume execution of the timeline at readily identified points. This capability will be supported both via ground intervention and via the on-board flight software. Examples of operations intervention to pause the timeline might include ground based determination that a solar flare is taking place and will interfere with the observations, or that the data recorder is full and must be dumped before the next observations are taken. In these cases, the flight operations team would use a real-time command to pause the execution of the timeline at a specified point. After the solar flare had passed, or the data recorder has been dumped, they would use another real-time command to resume operations.

There will be some types of activities that can naturally be paused and resumed (e.g. a sequence of identical exposures), while others cannot be so easily resumed (e.g. an individual exposure). The OPE event-handling mechanism will be designed to insert the pauses at the closest place in the observing sequence where resumption is possible. For health-and safety reasons, there will also be the option to terminate the current activity immediately without resuming.

The flight software system will provide a capability to modify the timeline after readily identified points. This capability will be used to support targets of opportunity, and might also be used if OTA performance has degraded and re-alignment is necessary. In this case, the timeline will be paused by a real-time command at a specified point, most likely at the end of a visit or complete observation, so that an alternate set of observations can be conducted. Operations will reload the command timeline from the point where it was paused, and then use another real-time command to resume execution of the timeline.

In many cases, the command to pause execution of the timeline will identify the point at which the timeline will be paused, and may be issued well in advance of the time that point will be reached. If the command to pause has not been executed, the command to resume execution will override that command. This means that it will be possible to reload the command timeline from a specified point without actually pausing execution of the timeline; the timeline would only pause execution if the specified point were reached before a command to resume execution was issued.

Examples when the on-board capability to pause the timeline would be useful include autonomous detection of filling of the data recorder by the flight computer, or detection of a solar flare via an on-board radiation monitor, or determination by the flight software that a momentum dump is required. In these cases the flight software will pause execution of the timeline, wait until the condition has cleared or necessary activities have executed, and then resume execution of the timeline. If necessary, a guide star reacquisition activity will be executed before the timeline is resumed. In some cases, such as a very high radiation background, the condition must be cleared by the flight operations team, and thus execution of the timeline would be resumed by ground command.

There will be cases (albeit rare) where it will be advantageous, from an efficiency perspective, for the flight system to recognize that an observation will fail and jump past that observation to the next one in the timeline. A clear case is where the flight system fails to find the expected guide star for a target. In this situation there is no point in carrying out the planned observations, so the system must have the ability to recognize the guide star failure and jump down in the timeline to the slew to the next target. Similarly, if some of the instruments have target acquisition modes, the failure of an acquisition to find the object is likely to be reason to jump past the planned science exposures.

On board radiation monitoring can be provided by a radiation sensor, or by monitoring data from the detectors. This monitoring might be combined with the fine guidance function, since cosmic ray detection may be required to avoid locking up on saturated pixels. Since solar flares and coronal mass ejections can last for hours to days, detection of high levels of radiation will result in suspension of the observation plan for extended periods. If the detectors are used to monitor radiation, they will continue to monitor radiation until levels drop and the observation plan resumes, or until levels exceed safe operation limits. If safe operation limits are exceeded, the controlling electronics will be turned off or disabled. In this case, they will have to be re-enabled by command from the ground, resulting in possible delays due to gaps between ground-station contacts.

3.5 Insertion of activities

There is likely to be a class of activities for which the flight software will automatically carry out activities not included in the preplanned timeline. These will be activities that cannot be predicted in advance, but the need for them can be detected on-board. Examples might include recovery from loss of fine guidance lock or insertion of off-nominal momentum unloading. The ultimate list of such activities will be determined much further in the design phase of the observatory. In such cases, the flight software will detect the event, pause activities at the next appropriate step in the execution of the timeline, carry out the recovery activity, return the observatory to the prior state, and resume the timeline.

The operations procedure for inserting a target of opportunity observation into the timeline will be the following. A command to pause execution of the observation plan after a specific observation will be uplinked. This will prevent the observation plan execution from proceeding into an observation that will be affected by the insertion of the target of opportunity observation. The new observation plan following that observation, containing the target of opportunity along with all other observations, will be uploaded and verified. Once verified, the new observation plan will be enabled (and the pause command, if not already executed, along with the replaced part of the old observation plan, will be deleted). Specifically, an observation will not be inserted into the observation plan. Rather, the latter portion of the observation plan will be replaced with a new version. This will eliminate the need for the ground system to ensure compatibility between the new and old versions of the observation plan.

The event driven nature of the execution of the timeline could result in cases where the timeline is ready to carry out an observation before it is desired or legal to do so. This might occur if an observation were scheduled to be taken just as the target visibility window opened, and the prior observation suffered a guide star acquisition failure. The system would need to detect this fact and delay the slew and observation until the target visibility window opened. In this case, it might be necessary for health and safety reasons for the vehicle to slew to a safe attitude until the next observation can begin execution.

3.6 Parallel Activities

There will generally be a need to carry out some activities in parallel, so the flight software and the timeline execution process will have to have that ability. At this early stage in the design the degree of required parallel activities, and the specific requirements are not yet known. Several examples can provide insight into different areas where this will be an important capability.

From an efficiency perspective, the time on target should be devoted to actual astronomical observations. The instruments may require some set up activities prior to the start of a new set of observations (e.g. turn on, filter/grating motions, etc), and some such activities between exposures on the same target (e.g. detector clearing, filter motions). Observatory efficiency can be maximized by having the on-board system execute these activities during slews and during the small-angle maneuvers used for dithering (section 4.10). Of more importance, there will be internal calibrations required to properly interpret the data. Some of these may be specific to the observations that are about to start. Others, such as exposures to measure detector dark current, may be needed on a regular basis simply to have up-to-date information on the state of the detector (e.g. hot pixels, latent images of bright objects, etc.). To the degree possible, we will want to schedule and execute these internal observations at times when the telescope cannot carry out external science observations, such as slew, settle, and momentum dump periods or execute the internal observation in parallel with an external observation using a different SI.

The software for requirements for parallel operations are probably not the major technical difficulty. The more likely difficulties will have to do with competing requirements for power, data transfer, and with the thermal effects of each activity on the overall thermal stability of the observatory. The need for parallel capabilities will have to be carefully considered since they will require extra care and expense in designing the flight and ground systems.

3.7 Observation Plan Executive

The Software Operations Working Group has begun to outline a software concept that would meet the general requirements described in sections 3.1 to 3.6. In this scheme, the overall flow of activities on NGST will be controlled by the Observation Plan Executive (OPE). The OPE executes an observation plan uplinked from the ground, but has the capability to

- initiate each activity specified in the observation plan based on preceding activities being completed and appropriate spacecraft subsystem states being achieved,
- insert into the activity flow any required house-keeping tasks that may be needed, and
- respond to certain pre-defined, non-nominal conditions that can interfere with the science timeline as they occur.

At the highest level, the observation plan will consist of a time-ordered series of visits. Visits themselves are relatively high-level constructs, specifying a logically complete set of operations that are intended to occur together in a fixed interval of time. An example of a visit might be a slew, followed by a guide star acquisition, followed by some spectroscopy together with some imaging in parallel. The detailed grouping and sequencing of these activities within the visit will be specified by the ground system, while the preconditions that must be met for the operations to begin will be determined either by the ground system or the on-board system. Specification of the exact start time for certain time-critical observations, for example, will be uplinked from the ground as part of the observing plan. Monitoring of the state of the momentum wheels and insertion of a thruster firing might be carried out autonomously by the OPE, but might also be restricted to within times or between activities specified by the ground system (for example, to ensure that those activities are executed during communications contact with the vehicle).

The principal job of the OPE will be to process the observation plan. The observation plan will specify a sequence of visits and a parallel sequence of spacecraft support activities for certain spacecraft support functions that do not control vehicle attitude. Vehicle attitude will be controlled only by visits. A time specification will be provided for each visit or spacecraft support activity, which specifies the earliest and latest time the visit or activity can begin execution and the latest time the visit or activity can end execution.

Each visit will contain a number of activities which execute in sequence or in parallel. Each visit or activity will specify conditions that must be satisfied before execution.

After the OPE has determined that all conditions have been met for initiating execution of an activity, the OPE will issue the appropriate command to the designated subsystem. Upon completion of the activity, the commanded subsystem will report back to the OPE the status of the command, i.e., whether it was completed with success or failure. The OPE will log the result and proceed with plan processing based on the status of the just-completed activity. If the activity succeeded or was optional, the OPE proceeds to processing the next activity. If the activity failed and was required, the OPE removes the remainder of the visit from the observation plan and proceeds to the next visit. If there is the possibility that partial completion of a visit will leave the spacecraft in a state from which it can not smoothly begin the next visit, the OPE will issue a standard clean-up command package to place the spacecraft in a standard inter-visit state. In addition to responding to execution requests issued by the OPE, each commandable subsystem will monitor its own state vis-à-vis the occurrence of various pre-defined non-nominal conditions and report such conditions to the OPE. There may also be subsystems with the sole function of monitoring spacecraft or environmental conditions about which the OPE needs to be informed. The OPE will respond to these reported conditions using a set of pre-defined rules that may depend upon the current state of the various subsystems. In general, the OPE will not stop execution of an activity. That is the responsibility of the Health and Safety function. The OPE may suspend execution of a subsequent activity (for example, if guide star lock is lost or particle flux exceeds a specified limit). Or the OPE may skip execution of subsequent activities (including an entire visit, for example if guide star acquisition fails).

In addition to such autonomous responses to spacecraft conditions, the OPE will be designed to allow easy real-time interaction with the ground, with support for appending observations to an existing plan, pausing or resuming the plan, inserting or deleting visits or spacecraft housekeeping activities, suspending at convenient points in the schedule, or interrupting activities immediately.

3.8 Onboard Data Management

In order to simplify overall operations the management of the on-board data storage system should be decoupled from the execution of the observational timeline. The flight software should automatically format and send instrument data to the on-board data storage system. There should be no need for the uplinked timeline to contain any indication of how large each data file is, or where it should be stored. This process is analogous to filling a silo from the top, continually dumping data into the on-board data storage system.

In parallel, the operations staff on the ground, or more likely an automated system on the ground, will establish the communications path to the on-board data system and will transfer data to the ground, releasing storage space as the data reach the ground. This is analogous to emptying the silo from the bottom. This process will be episodic, taking place only when there is an available communications link between NGST ground terminals and the NGST. The on-board data storage system should be sized to hold comfortably the amount of data that will be accumulated during the longest typical period of communications outage. The on-board data system should provide warnings to the timeline execution process in the event the recorder becomes filled, due to unexpected loss of communications. The execution of the timeline would then be paused until the data storage backlog was relieved and there was space available for new observations. There is no point in continuing observations without having a place to put the data.

The data recorder will provide capacity status to the observation plan executive. This will permit observation plan execution to be paused if sufficient storage is not available for an observation. It will also permit parallel observations to be skipped when available storage is below a specified threshold.

The use of a file transfer protocol provides a number of operational advantages. When science or engineering data is downlinked, a complete transfer and verification of a file will be performed before the file is released for deletion. If a file transfer is interrupted, it will be resumed on the next contact. This will be handled by the protocol, and will not require special capability of the ground system.

The data set generated in an instrument should have a unique set of identifiers for each individual exposure. The identifiers should be a part of the uplinked information in the timeline, and will provide the unique tag by which the data are known. Typically, this tag would include the proposal number, visit number, and exposure number within the visit. This information will be used in processing, distributing, and archiving the data. The data should be stored and accessed in both the on-board and ground systems via these unique identifiers. The operations staff should not need to go through a complex mapping scheme to know where in the recorder each observation is stored. From their perspective, selecting and downloading observation data sets to the ground should be seen as a file transfer process. Directory listings will be downlinked either automatically or on command. The standard operational mode will bring down the science data sets automatically in the order they were taken, but the operations staff will have the ability to interrupt this process and identify one or more specific observation data sets to be downlinked (perhaps by "clicking" on the data sets in a display of the on-board recorder and then "dragging" to the "ground"). Similar processes should be available for handling the transfer of engineering data from the recorder to the ground.

3.9 On-Board Data Processing

While the total integration times for typical NGST observations are long, the final images and spectra will be constructed from many shorter exposures, to allow cosmic-ray removal and dithering to remove instrument artifacts. During these short exposures, the detectors will be read multiple times at the start and end of integration in order to reduce readout noise.

The NGST yardstick communications concept consists of a single 8-hour contact period per day through an 11-meter X-band antenna. The data rate possible in such a concept is significantly lower than that required to transmit all of the individual detector readouts to the ground. Unless a larger data volume can be supported (by higher data rates and/or additional antenna), some form of on-board data processing will be required to reduce the data volume. Otherwise a large fraction of the NGST observing program will be limited by data rate rather than by the size of the primary mirror or the sensitivity of the detectors.

The simplest form of on-board processing will involve combining multiple reads to reduce readout noise. Loss-less compression to reduce the data volume will also be performed on-board. Both of these can be done robustly without degrading the scientific data and without posing a computational challenge.

We may also consider implementation of other compression algorithms or cosmic ray removal algorithms to further reduce the data volume if it is necessary to do so. The possibility of on-board cosmic-ray rejection is being studied (http://ngst.gsfc.nasa.gov/cgi-bin/iptsprodpage?Id-14; Bonati et al. NGST On-Board Data Management -- ESA Study). With perfect detectors the problem seems tractable without pushing the technological frontiers of on-board computing capabilities. However, departures from the ideal could lead to significant scientific losses due to non-optimal on-board processing, to scientific inefficiencies early in the NGST mission as the cosmic-ray rejection algorithms are tuned, or to significant costs as on-board software is rewritten to make the cosmic-ray rejection work in the light of detector peculiarities. The ESA study has pointed out, for example, that detector non-linearities may make it necessary to perform dark subtraction, flat-fielding, and non-linearity corrections prior to cosmic-ray removal. Prior experience dictates caution in this aspect of NGST design, as detectors on-orbit almost never behave exactly as expected. The primary concern about on-board processing is that, once it has been done, the original data are lost and cannot be reconstructed from the processed data that are downlinked. If on-board processing must be done (because increasing the downlink capability is simply too costly) it must be possible to uplink revised algorithms and, when needed, disable on-board processing and downlink raw data. Most importantly, great care must be taken to ensure that the on-board processing is correct and unbiased. Most scientific programs proposed for NGST require that the S/N of the data be close to that allowed by the background noise which means the cosmic ray rejection must work extremely well.

On-board processing will be necessary for target acquisition and wavefront control. Both will likely include some form of cosmic-ray rejection, but the requirements for that cosmic-ray rejection are much less stringent than for normal science observations because the exposure times are short and the measurements do not push to the S/N limits of the detectors.

3.10 Overview of Yardstick System Architecture

A schematic system architecture for the flight computer system was developed as part of the NASA yardstick design. The system represents current thoughts on how to fulfill the requirements discussed above with a minimum cost. Two themes will help reduce cost of development, test and maintenance of the data system: First, use Commercial Off the Shelf (COTS) hardware & software. Second: Commonality – use the same components, software and methodology as much as possible.

The onboard data system has two type of processing functions: Spacecraft related and Instrument support.

The spacecraft computer monitors the subsystems of the Spacecraft Support Module (SSM). The Integrated Science Instrument Module (ISIM) computer functions are:

- Communicate with SSM
- Observation plan executions (OPE), event driven
- SI mechanisms control
- FPA data acquisition
- Observation sampling algorithms

- Science Data compression
- Acquisition of Fine Guidance and FSM control

A preliminary description of flight software can be found at URL: http://ngst.gsfc.nasa.gov/public/unconfigured/doc_520_1/NGST-FSW-Req-I.pdf

An important aspect related to operations is the relationship between the SSM and ISIM systems. In normal operations (not safe mode) the ISIM computer will operate the NGST spacecraft. This computer will run the Operations Plan Executive (OPE), previously known as the Adaptive scheduler. The ground will upload the observation Plan, which the OPE will then execute by invoking the next sets of commands on NGST. More discussion on OPE can be found in section 3.7 of this document.

The proposed ISIM data system is based on the Power PC processor, which is estimated to meet the ISIM computational requirement. Several commercial real time operating systems exist. Standard Internet protocols could be used for inter-processor communication. Commercial bus architectures will be used where practical to reduce cost and to allow the use of commercial designs and test equipment. Commercial communication interfaces will be used where practical to allow standard off the shelf hardware to be used as simulators. This should simplify interface testing by allowing use of commercial interface documentation, test sets, training, and hardware.

3.10.1 Common Command and Data Handling

Details such as the number of processors or whether there is a "centralized" ISIM data system are under study. The Yardstick data system has two processors: SSM and ISIM. If there is more than one CPU, each processor will use the same operating system and Command and Data Handling (C&DH) components. The above approach is known as the "common C&DH". One solution for the common C&DH is the data bus architecture developed by GSFC is depicted below.

Figure 1

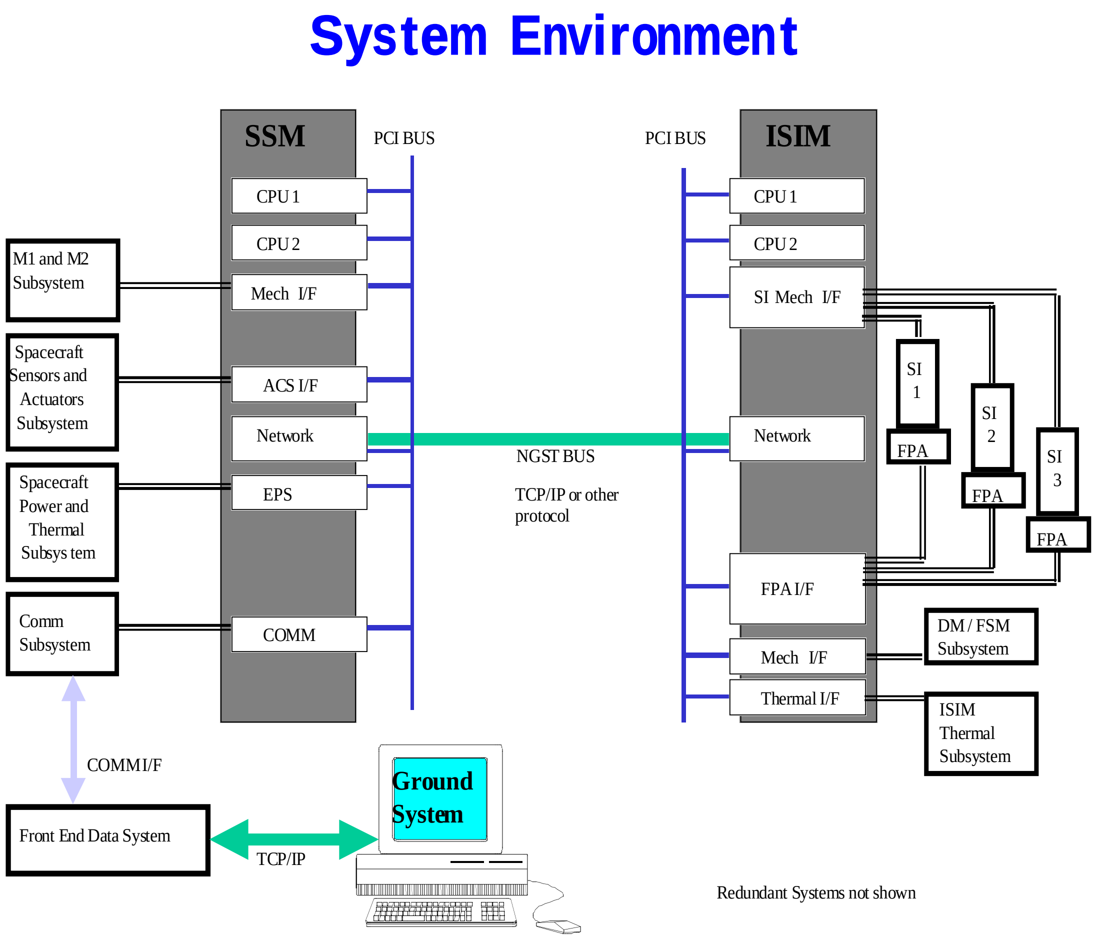

3.10.2 Data System hardware

The most challenging computing and memory requirements on NGST are those to support the NIR camera. The NIR camera produces data at a rate of up to 32 frames of 128 Megabytes every 1000 seconds. The data rate might increase by 4 to support cosmic ray removal, as discussed in the Communication Study (See References). The exact configuration and the total number of computers and CPU boards is being analyzed and will be decided by April 2000 (as part of ISIM cycle 1). A preliminary design of the ISIM data system can be found at URL: http://www701.gsfc.nasa.gov/isim/docs/aas/ISIM_data_system.pdf

4 Major Issues Affecting NGST Operations

4.1 OTA Stability and Alignment

The short-term and long-term stability of the OTA will figure prominently in the performance specifications for NGST. Major changes in the temperature distribution of the light shield can be expected after large slews, when the sun angle changes substantially. The on-board thermal management system can be expected to compensate for most such changes, but unless the temperature is actively and precisely controlled there will likely be residuals changes for the first several time constants of the system after a large slew, with consequent changes in the OTA and point-spread function (PSF) properties. Since some of the thermal time constants will be large, variations can extend for times comparable with typical observations. It is also likely that temperatures in the NGST structure will be continually changing at a low level, driven by pointing changes, temperature hysteresis in the materials, and other unknown effects. On longer periods, the end-to-end optical alignment of the telescope may also deteriorate, and will need to be verified regularly, probably on a time scale of weeks to months. The occasional required orbital adjustment maneuvers may vibrate the spacecraft. OTA variations will affect primarily the shape of the PSF; other optical instabilities may include variations in the flat field (especially in the vicinity of any field separators) and in the geometric distortion.

As the OTA thermal environment will change significantly with pitch angle, it may be necessary for the planning system to try to stabilize the thermal environment by scheduling observations to minimize changes in the absolute value of pitch angle outside of a certain nominal variation. It may also be necessary for the flight software to delay observations when either rapid temperature changes or direct measurement of the PSF indicate that the science observations would be significantly degraded. The thermal environment will also be monitored from the ground, and as on-orbit experience grows, scheduling constraints driven by OTA thermal stability will be refined.

Requirements will exist on both control and knowledge of the PSF. The shape of the PSF must be controlled to a level such as not to impair the NGST performance. Critical areas are the ability to resolve individual sources in crowded fields and the S/N that can be achieved for small, faint sources in the background-limited regime. PSF reconfiguration may not be completely automated onboard, thus the time scale for each reconfiguration operation will be a few days. Maintaining NGST efficiency then requires that the PSF remain stable, at the level required by performance considerations, on time scales much longer than a reconfiguration time scale, about a month.

A posteriori knowledge of the properties of the PSF at a level of a few percent pixel-to-pixel and of 1-2% in the enclosed-energy distribution, are needed for accurate photometry and PSF subtraction. Accurate PSF determinations may be possible on a weekly time scale, thus the PSF must remain stable at the required level over a week. At a minimum, the PSF properties must be measurable a posteriori as a function of both time and position in the field of view, by a combination of optical modeling and direct measurements. It is likely that temperatures in the NGST structure will be continually changing at a low level, since the thermal time constants will be long. There may be slip/stick effects, temperature hysteresis in the materials, and other unknown effects. It will thus be necessary to monitor the OTA performance by watching for changes in the point-spread function (PSF) in the science data.

It will be important to consider the scientific requirements on measuring and maintaining the figure and not just the engineering requirements involved in setting specifications. Scientists will want to beat the specification, if it is possible. Measurements at a higher level of accuracy than the specification will also be important to create predictive models of the alignment. Further study is needed to define how best to measure the PSF, what measurements to make on board and what to downlink to the ground.

The process to be used for OTA alignment still needs to be defined. It is likely to require a large set of alignment stars, since the slew and thermal settling times are large. We will need to understand and develop the observing, data analysis, and commanding processes to achieve alignment. For example, will we need to repeat the observations and analysis for verification?

We can expect that the alignment process will evolve substantially during the first year or two of operations. We should be able to automate and streamline the mechanics of taking and processing the data. We may be able to refine our operations model via improved limits on scheduling and better predictive models. We may also be able to tweak the on-board thermal control system to mitigate OTA changes. The timescale for these improvements will be months to years, since target selections will be driven by the science program and not by an engineering need to characterize the problem.

4.2 Momentum Management

NGST will use a large sunshield to shade the OTA and ISIM. The sunshield will be subject to solar radiation pressure, which will have a primary component normal to the sunshield and a smaller component tangential to the sunshield. The magnitude of the tangential component will depend upon the specular reflectivity of the sunshield and the inclination angle of the sun to the sunshield.

Solar radiation pressure will result in a force in the direction of the NGST center of mass, and, because the center of pressure on the sunshield is offset from the center of mass, a torque about the center of mass. The force in the direction of the center of mass will contribute to disruption of the quasi-stable orbit about L2 and add to the requirements for orbit maintenance. The torque about the center of mass will be compensated by reaction wheels. The reaction wheels can compensate for a certain amount of torque, after which an external force must be used to dump wheel momentum. The proposed method of dumping wheel momentum is through the use of thrusters. Any method available to NGST will result in an additional force in the direction of the center of mass, which will further contribute to disruption of the orbit about L2.

If reaction wheel momentum is dumped through the use of thrusters, we assume that this process cannot be executed during science observations, and may require additional time for spacecraft settling afterwards. Our concern is whether this will occur frequently, and whether the ground system must schedule observations to reduce the need to dump momentum in order to reduce the interruption of observations or the use of propellant, which is a consumable resource.

It is possible to design NGST in order to reduce the rate of momentum increase through an offset design of the sunshield or through the use of electrochromic panels. With an offset design, the sunshield is not flat, so changing pitch angle causes a shift in the center of pressure towards the center of mass. With electrochromic panels, reflectivity is adjusted at each end of the sunshield, also causing a shift in the center of pressure towards the center of mass. These designs reduce but do not eliminate the requirements on the reaction wheels and propulsion system. We assume for this document that these design elements are not incorporated in the NGST design, and we may modify our concept if they are. From an operations point of view it is highly desirable to have a torque-neutral attitude within the nominally allowable range of attitudes. Such an attitude may be desirable for an observatory safe mode, if nothing else.

The NGST goal is to provide reaction wheels of sufficient capacity that the largest pitch angle can be maintained for a minimum of 24 hours between momentum dumping by thruster operations. Analysis of this requirement (assuming a 4-wheel 45 degree pyramid configuration aligned with equal pitch component on each wheel; Isaacs 1999, in preparation) indicates a minimum reaction wheel capacity of 40 Newton-meter-seconds, although the yardstick design proposed the use of 20 N-m-s capacity wheels.

We assume that the NGST goal is met, and that momentum management will have the following impact on the ground system and operations:

1. We assume that momentum will usually be dumped before a major vehicle slew, in order to avoid interruption of the observation plan. The ground system will allocate time to dump momentum unless it is likely that a momentum dump is unnecessary. This decision will be based upon the pitch angle at initial and final attitude, the time spent at these attitudes, and the slew angle between these attitudes.

2. Because we anticipate that some NGST observations will require several days observing at the same attitude, it is possible that such observations will need to be interrupted to dump momentum. The ground system will anticipate the need to interrupt a long observation sequence to dump momentum, and allocate additional time to do so.

3. We assume that the event-driven observation plan execution will provide for the automatic insertion of momentum dumps by the flight software before slews and during long observations when needed. The observation plan will not in general explicitly request momentum dumps, either before a slew or during an observation sequence. However, it will be possible to explicitly request a momentum dump within the observation plan, in the event that we need a higher level of control. It will also be possible to modify the parameters that control the algorithm that controls when dumps are inserted.

4. The ground system will trend momentum as a function of pitch angle, as well as thermal data as a function of pitch angle and fuel use for momentum management and orbit maintenance. The trending of momentum and thermal data as a function of pitch angle will provide an indication of sunshield degradation and allow prediction of future fuel consumption.

5. Because momentum management will nominally require use of a consumable resource, the ground system should be able to reduce or even minimize momentum management requirements when necessary. This implies that the scheduling system will be able to limit pitch angles, while providing for exceptions when scientifically necessary. The momentum neutral attitude may be at a non-zero pitch angle (in fact, the yardstick does not have a torque-free attitude within the nominal range of pitch angles).

6. Additional capabilities to be considered include imposing a scheduling bias to minimize momentum dumps by scheduling observations with preference to times when the torque is minimized. The yardstick design does not have a torque-free attitude within the nominal range of pitch angles, because the center of mass is offset from the normal to the sunshield at the center of pressure. This is an unfortunate design element. If the center of mass were aligned with the center of pressure, then momentum dumps could be minimized by scheduling observations at alternating pitch angles. This scheduling approach would also be consistent with a scheduling approach that minimized thermal variation (which changes with the magnitude of the pitch angle).

7. There are a number of failure modes that must be considered. These include reaction wheel failure, unexpected use or loss of propellant, and premature sunshield degradation. Each of these conditions will require the ground system to implement immediate restrictions on scheduling to reduce momentum. We anticipate that the most likely response will be to impose an immediate limitation on the pitch angle, while the long-term response will be to implement scheduling restrictions that balance momentum by shortening individual observations and scheduling at alternating pitch angles.

If the design requires momentum dumps more frequently than about once per day it will affect mission lifetime, observing efficiency, and operations complexity. Momentum management thus has the potential to have a serious impact on operations, but an appropriately designed set of reaction wheels should mitigate the problem. A system requiring momentum dumps once every 24 hours, with short overhead required for the dumps, will not be a significant scheduling or operations problem. Nonetheless in this optimal case, some amount of scheduling, modeling and control and engineering trending will be required of the ground system.

Further Analysis:

- Approximate the time for uncompensated torque to rotate the spacecraft into an illegal or unsafe attitude.
- Estimate slew times based upon a reaction wheel capacity that meets operational goals.

4.2.1 Design Considerations

If no mitigating design features are incorporated into the sunshield, the reaction wheels must be larger than the 20 N-m-s capacity wheels originally anticipated. Another factor in reaction wheel sizing is vibration, especially at low frequencies. To avoid excitation of low frequency modes, the reaction wheels must have isolation mounts, and also must be run at speeds above a certain minimum (for example, to avoid excitation of a 12 Hz vibration mode, the lower limit on wheel speed is 720 rpm). This restriction reduces the usable speed range for the reaction wheels, and thus shortens the time between momentum dumps. One advantage of larger reaction wheels is a reduction in slew times, which will improve observing efficiency.

The reaction wheel capacity to support ±25 degree pitch angles for up to 24 hours in a 4-wheel 45 degree pyramid configuration is 37 N-m-s. This capacity would allow observations up to 86 hours between momentum dumps at the yardstick minimum momentum attitude.

One candidate for this wheel is the Ithaca-E, which has a 50 N-m-s capacity (and a 0.3 N-m torque capacity). With these wheels, the range of time between momentum dumps would be 32 to 117 hours.

4.3 Ground stations and operations

The NGST yardstick provides one high data rate ground station for high rate science and engineering downlink and low rate command and telemetry uplink/downlink. We have recommended that a second low rate ground station be provided at a complementary location for initial operations and for contingency purposes. For budget reasons, this second station may only be available for contingency purposes.

As we discuss later, it might be worth providing a second high-rate ground station for both command and telemetry and science downlink contingencies. If the primary high-rate ground station is located at high latitude, then during summer months the contact duration will be short. This will increase the on-board storage requirements. During solar maximum, frame rates may be higher to mitigate cosmic ray impacts, resulting in higher data volume. Using a second high-rate ground station during summer months of solar maximum will alleviate data volume problems.

If the high-rate ground station is located near the longitude of the NGST Operations Center, communications contacts will occur during the middle of the night. During initial operations 24-hour staffing will thus be required to monitor spacecraft systems as well as evaluate science data. However, we plan to quickly reduce or eliminate staff during this shift if NGST operations do not encounter problems that require manual intervention.

Communications with NGST will be automatic. Contact will be established and real-time engineering telemetry will be captured by the ground system. Flight software status and event logs will then be downlinked and processed by the ground system to determine whether problems were encountered during observation plan execution while the vehicle was out of contact with the ground. If problems were encountered, and if the severity of the problems is high, then the ground system will automatically contact on-call operations personnel.

Observation execution times will be extracted from the status and event logs and used to update the observation plan. The updated observation plan will be uplinked for the next day(s) observations. Several days worth of observation plan will be maintained on-board in case a ground contact is missed. Additional support files will also be uplinked as needed.

Recorded engineering telemetry data will then be downlinked and made available for analysis. Finally, recorded science data will be dowlinked until the end of the contact period is reached. Engineering data will be recorded in files covering set time periods, and science data will be recorded in files for each observation. As each file is received, decompressed and verified, it will be released for deletion on-board. Critical files may be identified for manual release. In this case, these files will not be deleted on board until they have been reviewed by operations personnel.

4.4 Data Volume and Downlink Capacity

The science data volume from NGST will be large for two reasons. The detectors are very large format, and the radiation background from galactic cosmic rays and solar proton events is high and requires frequent readout of the detectors. We assume specific characteristics of the radiation background, impact on the detectors, and an approach to processing cosmic rays.

The radiation background consists of galactic cosmic rays and solar proton events. Galactic cosmic rays provide a constant, isotropic flux of 5 particles per square centimeter per second. Solar proton events are variable throughout the year and over the 11-year solar cycle. During solar minimum, solar proton events provide an isotropic flux of less than 2 protons/cm^2/sec through 90% of the year. During solar maximum, solar proton events provide an isotropic flux of less than 5-10 protons/cm^2/sec through 90% of the year (we have used 5 for our analysis).

Radiation testing of InSb detectors for SIRTF showed at least 5 pixels were impacted per cosmic ray event. In this case, one or more pixels were hit directly by a cosmic ray, at least one of which saturated and produced charge overflow into the four adjacent pixels. Saturated pixels exhibited persistence effects for some time after the detector was cleared. These results are valid for low energy cosmic rays, against which we expect the detectors to be shielded. Analysis of NICMOS on-orbit data shows that 1.6-2.0 pixels were impacted per cosmic ray event. If we scale these results from 40 micron pixels of NICMOS to 27 micron pixels for NGST, we estimate an impact rate of 2.0-2.5 pixels per cosmic ray event. For this analysis, we have assumed that less than 10% of pixels are saturated by lower energy cosmic rays, and that an average of 2.5 pixels are impacted per cosmic ray event.We assume the following approach to processing cosmic rays. Suppose an exposure time of 10,000 seconds is required, and the detectors must be read and cleared every 1000 seconds (a sub-exposure) in order to avoid saturation and cosmic-ray degradation on the one hand and minimize read-noise effects on the other. We assume that multiple reads of the detector are averaged over a short interval to reduce read-noise. Between these sub-exposures, the image may be dithered to improve spatial resolution and compensate for hot pixels bad regions on the detector and sensitivity variations. We add one or more sub-exposures to the total required exposure time to provide additional exposure time which is used when cosmic rays are detected and the exposure time on the effected pixels is removed. We produce a zero frame at the start of each sub-exposure, to establish the bias levels and to help identify persistence effects, and we produce one frame at the end of the sub-exposure. We may also produce additional frames during the sub-exposure, which improves the removal of cosmic rays but increases the data volume.

The trade we study is that of overhead (number of additional sub-exposures) versus data volume (based on number of frames per sub-exposure). The controlling parameter is the percentage of pixels that achieve the required exposure time after processing to remove cosmic rays. The following graph shows the result of this analysis. We plot the percentage of pixels that achieve the required exposure time, for an overhead of 5%, 10% or 15% additional exposure time, versus the number of frames per sub-exposure (excluding the zero-frame mentioned above). We provide this data for the radiation background at solar minimum and at solar maximum. We also plot the data volume, which varies by the number of frames per sub-exposure, for cases where we use just the NIR camera (no parallels) or all instruments (parallels) and for the cases where we can achieve 2:1 or 3:1 data compressions. The data volume estimates assume 85% efficiency (time during which exposures are being taken, versus slews, acquisitions, etc.), continuous engineering data, and 15% communications overhead.

Figure 2

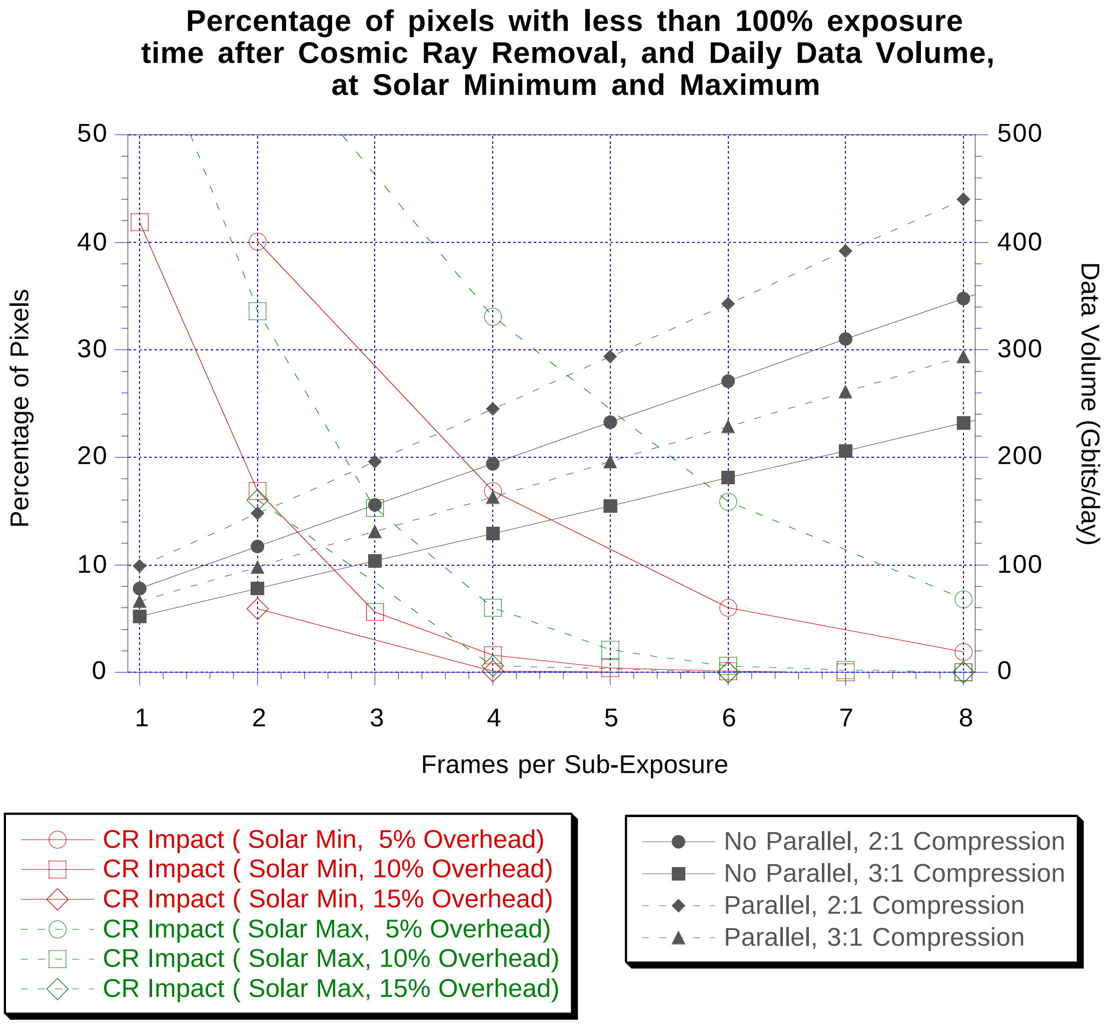

As an example, suppose we want 90% of pixels to have the required exposure time, and we are only willing to accept 10% overhead (one sub-exposure). Then during solar minimum we would need to take 3 frames per sub-exposure, and during solar maximum we would need to take 4 frames per sub-exposure. If we schedule full parallel exposures and can only achieve 2:1

compression, then during solar minimum we would generate about 200 Gbits of data per day, and during solar maximum we would generate about 250 Gbits of data per day.

The Canadian Space Agency has proposed a 46-meter antenna for use by NGST. This antenna would provide a minimum of 700 Gbits/day, which is sufficient for all cases illustrated in the plot. We highly recommend a comparable solution, and we recommend sufficient on-board data storage to store at least two-thirds of the maximum data volume, based upon the final determination of the degree of data compression that can be attained and whether full parallel operations will be supported. Our prudent recommendation is 135 Gbits of on-board storage, which would support non-parallel operations with 2:1 data compression during solar maximum. In this case, parallel operations would be restricted during summer months and expanded during summer months when the ground station has longer contacts (and thus the demands for on-board storage are reduced).

Science proposals may include options to specify the percentage of pixels that will achieve the required exposure time, allowing scheduling tools to set the number of frames per sub-exposure and number of additional sub-exposures taken to balance data volume and overhead requirements.

The scheduling system will ensure that on-board storage capacity is not exceeded, but if sufficient capacity is provided, this will not be a difficult restriction to enforce. If sufficient on-board storage capacity or data transfer capacity is not provided, then this will result in a decrease in scheduling efficiency and an increase in overhead. In this case, the scheduling system might have to adjust overhead and frame rate depending upon schedule conditions. This is a complexity we would prefer to avoid.

Further discussion of communications system tradeoffs can be found at http://presto.stsci.edu/vision/vision_se/ngstOpsConceptWG/CommunicationsStudy/index.html.

4.5 Slew Characteristics

Slew duration will depend upon the capacity of the reaction wheels. Higher capacity reaction wheels will not only reduce slew duration, but also the time required between momentum dumps. Reaction wheels that support momentum dumps no less than once per day will support 90 degree slews at less than 30 minutes (including settling time). Settling time after a slew will depend upon the rigidity of structural joints; rigid joints will result in smaller settling times but may result in mechanical disturbances during thermal settling. A settling time of a few minutes is assumed to be consistent with structural joints that do not exhibit stick-slip problems during thermal settling.

The sun is required to be within 5 degrees of the X-Z plane. A typical eigen-axis slew from one attitude to another would violate this requirement unless the eigen-axis were within 5 degrees of the sun or Y-axis. To maintain this requirement, the vehicle slew can be separated into two eigen-axis slews, one which adjusts roll and pitch angles, and one which rotates about the sun vector. We recommend executing the roll/pitch slew first, as this would minimize the thermal impact of the slew as well as maximize the thermal settling, which would be generally unaffected by the slew about the sun vector. We will allow for a star tracker acquisition after the roll/pitch slew if that maneuver is large.

Small angle maneuvers, of order up to 1-2 FOVs, will commonly be required. These maneuvers will be essentially jitter free, and will not require settling time. Multiple clears of the detectors will be executed during these maneuvers (to reduce persistence of cosmic rays and bright objects).

4.6 Solar Flares

The cosmic ray activity pattern for NGST will be very different from that in a low Earth orbit. The frequent, large, but predictable swings due to the South Atlantic Anomaly will not be present, thereby removing the related orbit-based scheduling constraints. On the other hand, the spacecraft will be much more exposed to direct hits of particles from the Sun, which will cause irregular fluctuations in cosmic ray activity. Currently we estimate a mean flux of 5 particles $cm^{-2}s^{-1}$, which translates into of order of 10% of the pixels affected by cosmic rays for a 1000s exposure (27 micron pixels, 2.5 pixels affected per event, 50 mil Al shielding). Thus cosmic rays may well be the limiting factor for exposure length on low-background data, and the optimal exposure length may vary somewhat depending on observing strategies and detector characteristics. It may be worth investigating the cost and benefits of an adaptive exposure manager, whereby the rate of cosmic ray hits is monitored and the exposure ended when appropriate. Since the telescope is exposed directly to the flux of low-energy particles from the Sun, shielding will have a significant effect on cosmic ray hits, and cosmic ray activity will be difficult to predict accurately before launch.

A more serious operational issue will be solar flares and coronal mass ejections, which occur infrequently but can increase the total cosmic ray activity by up to 5 orders of magnitude. Obviously observing is out of the question during a flare. A more important question is whether we will need to shut off some components to protect them from damage, either instantaneous or cumulative over the course of the mission. If instantaneous damage is the danger, we also must consider whether we need advance warning so that the telescope can be shut off before the peak flux reaches it requiring continuous communication contact, or we can rely on onboard sensors? Advance warning may come from other satellites inside the Earth's orbit, and will need to be prompt and reliable; the spacecraft will need to react to the warning quickly and automatically. We will learn more about the particle environment at L2 from the forthcoming MAP mission. Clearly, experience in the first year or two will be important here.

4.7 Science Instrument Operations

The operation of the complex multiple science instruments of NGST will likely require balancing optimal efficiency and breadth of capabilities against operational complexity. Some details will depend on the specific instruments selected, but general considerations can be made.

It will be desirable to command many aspects of the SIs in parallel (setups etc.) even if the functionality for parallel observations or data taking is not implemented. Ideally, operation of one SI should not affect the others. Operations become much more complex if the ability to observe with or perform housekeeping functions on one instrument depend on the state of a different instrument. However, the reality may be that power restrictions or limitations on thermal output place restrictions on parallel operations. The flight system should be designed to allow these restrictions to be incorporated into the general event-driven philosophy. For example, if the thermal conditions necessary to allow a parallel operation to proceed are not obtained, the operation may be skipped.

It is highly desirable from an operational point of view to keep the detectors at a constant operating temperature, to avoid long overheads for warm up or cool down sequences. If such sequences are necessary and involve large overheads, that will affect the overall philosophy for planning observations.

4.7.1 Parallel Operations

In many current NGST concepts, several science instruments occupy the focal plane simultaneously, so that “true” parallel observations (separate instruments observing different regions of the sky at the same time) are possible in principle. Parallel observations increase the complexity of operations:

a) telescope operations driven by the primary instrument, such as dithering schemes, may affect the observations by other instruments; an appropriate decision tree must be established in the scheduling process, including the on-board adaptive scheduling system if implemented;
b) conflicting requirements must be identified and avoided;
c) some science instruments may need special operating modes to carry out science in parallel (e.g., point-and-shoot for multi-object spectrograph);
d) power and control management increases in complexity;
e) data acquisition and processing rates are more demanding;
f) stringent requirements must be imposed on cross-instrument interference (induced vibrations, baffling of calibration lamps, and so on).

On the other hand, parallel operations may increase substantially the productivity of the observatory, especially if some of the major planned science programs can be advantageously executed in parallel. Furthermore, overlapping SI capabilities (e.g., imaging mode of multi-object spectrograph) may enable some programs to be executed more efficiently in a pseudo-parallel mode, by simply increasing the effective field of observation. Whether the additional cost and complexity of true parallel operations will be justified depends on the perceived scientific worth of parallel capabilities and the actual instrument designs.

If true parallels are not deemed justified, a more limited capability for “internal” parallels (other instruments carry out internal, non-pointed observations while the primary instrument observes sky targets) is worthwhile. Without natural observation interruptions (occultations), there will be fewer opportunities to carry out the internal calibration observations (darks, flats, lamp wavecals, etc) which on HST represent the majority of the on-board calibration time. The ability to execute such observations in parallel to other instruments' science observations will ensure that less science time will be lost to calibrations, improving the observatory efficiency, and that internal calibrations and monitoring of the instruments' stability can be carried out on a regular basis without competing with science observations. On HST, the generous allocation of calibration monitoring time, made possible by the availability of internal calibrations, has improved our knowledge of the instruments and allowed early detection of problems, such as the position shifts of WFPC2 chips. Internal parallels are largely unaffected by considerations a) through c) above, although they will affect the power and data processing management and the possible interference (d) through f)).

Some degree of parallelism between spacecraft and SI activities will also be needed to ensure efficient execution of observations. For example, any instrument setup or reconfiguration activities prior to a new observation (turn on, motions of filters and gratings) should be scheduled by the on-board system to execute during the spacecraft slew. Some of the internal calibration observations may also be carried out during slews, thus improving the efficiency of the telescope even without true parallel capabilities.

4.7.2 Flexibility vs. Complexity

Instruments capable of diverse observations can offer significant scientific advantages. A flexible instrument can offer better tailoring of observing parameters (choice of filter, choice of grating, observing mode for a multi-object spectrograph, and so on) which can enhance its scientific return. In some cases, multifunction instruments offer overlapping capabilities, and therefore some protection against failure - as long as their replacement modes work well.

On the other hand, the cost of building and operating an instrument increases with its complexity. More functionality implies more complex commanding software, more parameters to be considered in scheduling, and more time spent in taking and analyzing calibration data. Different types of complexity will impact each element in a different way - software requirements may not increase much with the number of spectral elements, but calibration time will. Rarely used modes can skimp on calibration, but may well require as much commanding work (and thus cost) as more commonly used modes. High precision target acquisition (aperture spectrography, coronography) requires two-way interaction between the instrument and the pointing control system. Adjustable-mask multi-object spectroscopy may require substantial on-board processing of the acquisition image, without necessarily iterating with pointing control.

A flexible instrument is only justified if 1) it can offer significant quantifiable advantages in scientific return, and 2) the advantages offset the increased work, and possible decreased efficiency, associated with instrument complexity. Multiple functionality can offer contingency against failure of another instrument, but the replacement capabilities must be evaluated critically to ensure that they are adequate.

4.7.3 Calibration

Although the specifics of instrumental calibration will depend on the detectors and instruments chosen, some general requirements can be envisioned for both imagers and spectrographs.

4.7.3.1 Imager

Areas that require calibration will generally include:

- Gain and read noise (needed separately for the noise model)
- Analog-to-digital converter
- Bias and bias stability
- Internal background and dark current
- Filter-dependent telescope background
- Sky background
- Flat field
- Photometry (characterization of filter throughput, including detector QE)
- PSF
- Geometric solution

The various sources of noise and background all behave in different ways, and need to be calibrated differently. Analog-to-digital converter errors are likely to be extremely stable, and only need occasional recalibration. Bias, as long as it remains stable over typical observing times, is removed in the normal observation sequence. The methodology of monitoring and maintaining will be determined by the camera design. Similarly, dark current is likely to be stable over short time scales. A full sequence of combined bias and dark observations will probably need to be repeated on a regular basis, for example to monitor any long-term dark current degradation due to particle damage and other aging effects. Both internal and telescope background can have structure on all scales, but remain approximately constant for small slews, at least as long as bright sources (such as the moon) are properly baffled. Dithering strategies may be needed to correct for this type of background, especially in the thermal infrared. True sky background may be nearly constant over the field of view, but it can vary substantially with large slews. Accurate determination of the flat field requires the combination of pixel-to-pixel variations, probably using internal light sources, with large-scale structure which can be determined from extensive sky background observations. Any spectral dependence of the flat field will also need to be characterized with internal lamps. Photometric throughput will require observation and monitoring of known stable sources several times per year. With the pointing restrictions of NGST, this may require establishing a network of sources that can be related to one another over time.

A sufficient calibration program with the imager requires:

- Availability of stable internal calibration light sources capable of producing adequate continuum signal throughout the instrumental observing range. Two (or more) light sources with different spectral energy distributions may be needed to solve for wavelength-dependent pixel sensitivity.
- An extensive program of monitoring via internals, including a full sequence of bias/dark combinations, internal background, and internal (small-scale) flat fields. The latter will be needed separately for each spectral element.
- Regular external observations to characterize and monitor large-scale properties of the flat field, photometric throughput and stability, PSF and properties, and geometric solution. Photometry and PSF will be needed at regular intervals for each spectral element, while large-scale flat field and geometric solution can probably be maintained with observations in a subset of filters.

Internal observations will benefit greatly, in terms of telescope efficiency, from the ability to carry out observations either during slews or as parallels.

4.7.3.2 Slit Spectrograph

A slit spectrograph will need additional calibrations besides those required for an imager. Some examples are:

- Wavelength zero point calibration
- Wavelength solution
- Slit throughput
- Slit alignment
- Line shape calibration

The wavelength solution and zero-point define the relationship between pixel addresses and wavelength, with the zero-point corresponding to a simple shift with no change in scale or distortion. Both calibrations can probably be achieved with internal light sources; multiple lines will be needed for the wavelength solution. Wavelength zero point calibrations may need to be repeated for each observation. Slit calibration will require external observations to determine the slit attenuation function, as well as internal observations with a continuum lamp in order to locate the slit with respect to the focal plane. Line shape calibration will probably require external observations.

Operationally, slit spectrograph calibrations will need an adequate set of calibration light sources, both line and continuum, as well as appropriate instrumental modes to execute the observations. Many of the calibrations will be internal and require moderate instrumental stability, but throughput and line shape will require external observations of suitable sources.

4.7.3.3 Multi-Object and Integral-Field Spectrographs

The possibility of multi-object spectrographs, either with integral field or with configurable apertures, introduces the need for additional calibrations, especially in the response properties of individual apertures and in their individual wavelength zero points. The details will depend on the construction of the instrument. In general, the calibration process may be somewhat easier with the fixed apertures of an integral field system, although the relative position of each source within its aperture may be a difficulty for high-precision wavelength calibration if the aperture is larger than the source size (including PSF). Wholesale wavelength and throughput calibration of a large number of apertures can be made easier with line sources. For configurable apertures, filters will be needed to enable simultaneous calibration of a large number of apertures without confusion between different spectral lines. Because of the potential difficulty of calibrating thousands of apertures in flight, the fundamental relative calibrations must be performed during I&T or at the piece level. These will be verified in flight.

4.8 Guide Star Acquisitions and Tracking

The current major issue regarding the acquisition and tracking system is whether it will derive from the science camera, or be a separate system. In the Yardstick design, the guiding system uses part of the primary NIR imaging field of view as input to the pointing control system. In this mode, a guide star would be identified for each pointing, using either a catalog or preliminary imaging. A small region of the detector around the guide star would be read out very frequently - about 30 times per second. The instantaneous information on the position of the star in the detector would then be used to drive the fast steering mirror and maintain telescope pointing stability.

A major motivation for having the guider be part of the camera is the fact that the cost per pixel is currently the biggest limitation on the size of the NIR field of view (FOV) of the science camera. With this assumption, some portion of the NIR field of view must be sacrificed to the guiding system regardless of whether or not a guiding system separate from the NIR camera is chosen. A guiding system that employs only a small portion of NIR FOV around the guide star would presumably result in the minimum loss of FOV. Furthermore the stability of the guider-to-camera geometry is much easier to achieve if the guider is part of the camera. Indeed, guiding with the fast-steering mirror may introduce significant field distortions. Having the guide stars as close as possible to the science focal plane will minimize the impact of such distortions.

The additional obvious advantage to using the NIR camera for the guiding system is that one less instrument needs to be developed. On the other hand, if geometric distortion, stability, cost, and reliability issues could be resolved, a dedicated guider offset on a mechanical x-y stage, would offer the potential of deploying a small number of pixels over a much wider area of sky. A separate guider would require accurate calibration of the field distortion and the distortion stability. Stability and accurate calibration of the camera distortion will be required whether or not it functions as the guider; however if it is to be the guider for the other science instruments, the requirements on knowledge of its position and orientation relative to the other instruments are the same as for a dedicated guider. In other words, using the camera as a guider does not lessen the requirement for a stable and well-calibrated focal-plane geometry. A dedicated guider could have a different plate scale, operate further into the optical where the PSF is sharper, and use filters optimized for guiding. The full square field of view of the NIR camera could be used for science observations, and concerns about losing the guide star as the telescope orientation changes or as the observations are dithered to fill in the detector gaps would be mitigated. Provided a larger field of view can be searched for guide stars, operations with a dedicated guider will be simpler than with the NIR camera functioning as the guider.

The Yardstick ISIM envisions a NIR FOV of ~16 sq arcmin, a small area in which to identify potential guide stars. (By contrast, the HST FGSs have areas of 69 sq arcmin each.) There are currently no complete star catalogs for the NIR, although the 2MASS survey will be deep enough to generate one. Estimates of NIR star counts suggest that on average there will be one star per field bright enough for guiding. Specifically, depending on the star-count model used, it is estimated that 95% of the fields near the Galactic Poles will have at least one star brighter than K = 15.5 to 16 (K_{AB} = 17.5-18) (Spagna 1999; http://ngst.gsfc.nasa.gov/public/unconfigured/doc_458_2/IRCounts.pdf). This suggests that positions of guide stars relative to targets will be an observing constraint for at least some programs. A NIR guiding system would have to deal with the eventuality of the only guide star being in the same camera as the target. Increased operational complexity would result from having to determine which camera to use for guiding.

Guiding requires a broad-band filter to increase the number of photons detected, thus reducing the intrinsic pointing noise. For a four-camera NIR imaging instrument, this means that–unless the primary observation also requires a suitable broad-band filter–a full camera (25% of the data) is lost to the science observations. On the other hand, many of the narrowband observations envisioned for NGST either do not suffer greatly from a reduced field of view (Satyapal 1999; http://www.ngst.stsci.edu/studies/ study99_narrow_band_guide). Of course, the optical design must support observing with a narrow-band filter while guiding with a broad-band filter.

Dithered observations (see 4.10) would also complicate the guiding process. Each pointing would have to be broken into several exposures with slightly different offsets. These motions may be a large fraction of the size of an individual detector, forcing the selected guide star to move from detector to detector in the course of an imaging sequence. In addition, the process of minimizing the impact of the interchip gaps will increase the chances of the guide star falling in the gap for at least one of the dithered observations.

Issues that need further study as part of the guider selection:

- Can guiding with the NIR camera really only sacrifice a small area around the guide star?
- Is more than one guide star required per observation (e.g. to allow for dithering)?
- Based on NIR star counts, what is the probability of having a guide star for different plausible observing scenarios (e.g. observations taken over a range of orientations with dithering to fill in the detector gaps; or spectroscopic and mid-IR observations of the same field observed by the camera).
- Which would be more cost- and science-effective, development of a NIR catalog, preliminary imaging, or autonomous guide-star selection? Which would be simplest operationally?
- What is the failure rate likely to be due to “natural” causes, such as binaries, bad magnitudes, etc.?

4.9 Observation Strategies - Roll Flexibility

Note that throughout this section, the **nominal roll** is defined as the position angle of the vector from the target to the Sun, projected on to the sky as seen from NGST at the time of observation. This is evaluated in degrees, measured east from **ecliptic** north. Thus for targets in the ecliptic plane, nominal roll is always ±90°, while at the ecliptic pole, nominal roll goes through all values at approximately one degree per day.

Many astronomical observations desire a particular telescope orientation or the ability to observe the same area of the sky without rotation of the field of view for long periods of time. Multi-object spectral observations may desire a particular orientation as well, to avoid confusion in the dispersed light. HST observations are constrained to be made within 30° of the nominal roll angle. To remain cold, the NGST telescope assembly must remain in the shadow of the sun shield. Thus, the NGST roll will probably be much more constrained. For example, at least in one incarnation (Bely 1998), the Yardstick Design only allows a maximum ±2° off nominal yaw of the sun shield which translates into about a ±2° off nominal roll about the boresight. A wider sun shield would allow a greater off nominal roll range but would also create larger radiation induced torques and increase the complexity of mechanical issues (deployment, instabilities, etc.). A balance will need to be made between the size of the sun shield and the limitations it places on observations.

The limitation in off-nominal roll implies that for NGST targets within about 30° of the ecliptic plane (50% of the sky), only orientation angles within a few degrees of +90° or –90° will be possible. Overall, the available roll for a space telescope is a function of the target's ecliptic latitude, the allowed off nominal roll, and the allowed target-sun angle. Using the constraints of the NGST yardstick design, a comparison of the available roll angles for NGST and HST as a function of the sine of ecliptic latitude is shown in Fig. 3. The sine of the ecliptic latitude is used in the plot as it represents the cumulative fraction of the total sky area as measured from the ecliptic plane.

Figure 3

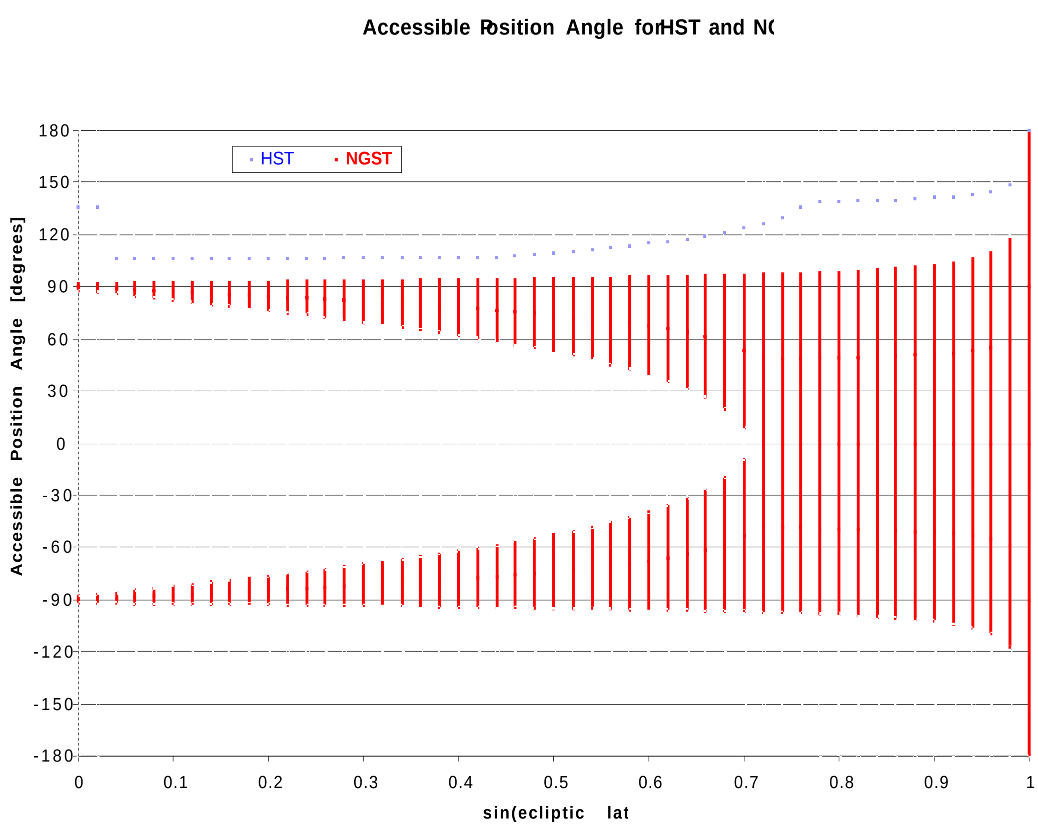

The permissible ranges of motion are different for the different sunshield axes in the Yardstick design, and in most likely NGST designs. Thus the scheduling system will have to take into account the sunshield geometry and compute the allowable boresight rolls for the desired targets as a part of the computing the observing windows. Because many of the likely NGST targets are at high ecliptic latitude, providing greater latitude in boresight roll in the NGST design may be as least as advantageous as increasing the field of regard.

The rotation of the nominal position angle over time will also restrict the duration of an observation that requires or desires a fixed orientation. For example, with ±2° off-nominal roll capability, a single observation using a fixed orientation of a target near an ecliptic pole will have a maximum 4 day duration. A single, fixed orientation observation in the ecliptic plane will be limited by the solar exclusion constraint to a duration of about 50 days per year. The maximum allowed duration for a continuous, optimally oriented, fixed orientation observation as a function of the sine of the ecliptic latitude and the allowed off nominal roll is shown in Fig 4.

Figure 4

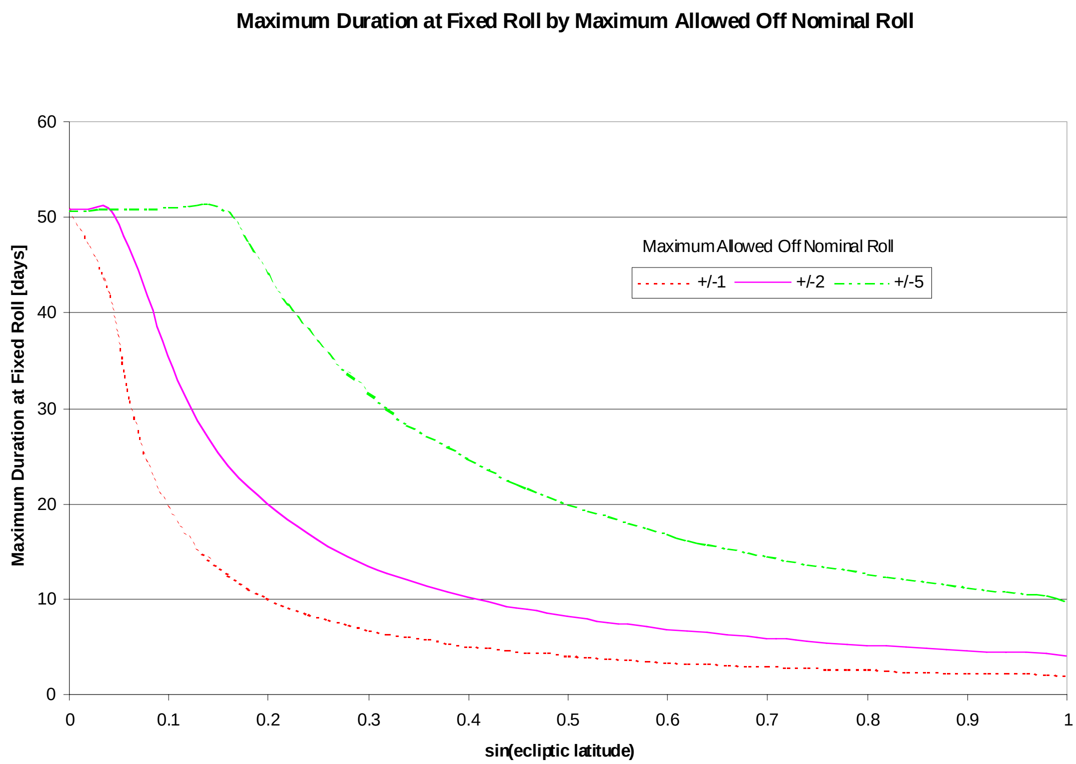

Except for short, simple point and shoot observations lasting only a day or two most observations will be affected by the orientation constraint. To determine the overall impact of this constraint, we will need to answer the following questions: Does the orientation restriction directly affect the science or is it an efficiency impact? How many observations of a particular type will we expect to execute? For all types of observations, what will be the best observing strategy?

Roll will be an issue for long observations, especially mosaics. If NGST has a multi-object spectrograph, then roll restrictions will be determined in order to optimize the number of targets that can be observed without interference with other objects. Roll restrictions will be imposed for observations with limited guide star availability. Finally, roll restrictions will be imposed for repeat visits. Some roll flexibility will be provided to the proposer, while the remaining flexibility in roll will be reserved for the planning system in order to provide scheduling flexibility to the ground system. If the spacecraft-imposed roll restrictions are very tight, event-driven scheduling will become difficult, since the roll angle will depend critically on the time of the observation.

4.9.1 Mapping (Mosaics)

Some of the DRM involves mapping an even larger area than the 4 arcmin square planned for the NIR camera. Such maps will comprise separate images at different pointings. These images will presumably be the product of a series of exposures dithered (see Section 4.10) around a given target pointing. The exposures of a given target pointing could be combined on board the spacecraft to limit the data rate. The mosaicing of these combined images to make the larger map would be done on the ground because the different target pointings may be made at significantly different times and hence, orientation angles. The highest observing efficiency will occur if the separate targets composing the map are made at nearly the same position angles or additive multiples of 90°. This would allow the creation of a mosaic that minimized unwanted overlapping of adjacent fields. Over most of the sky, using this technique to maximize the usable time on target will restrict the observations' position angles close to –90°, +90° and at mid-latitudes, also 0°. In addition, for about 80% of the sky the maximum usable time will be about 20 days per year with the individual observations having no scheduling flexibility. The exception to the above rule would be for targets within about 10° of the ecliptic, where about 100 days per year would be available, half at +90° and half at –90°.

Two of the proposed observations in the DRM will each require approximately 60 days of observations to map two small areas of the sky at high ecliptic latitudes. Using the above observing strategy would imply that these programs would take over 3 years to complete. To complete the observations in less time implies that the orientation of the observations cannot remain fixed.

However, rotation of the individual fields in a mosaic can cause holes and overlapping images to occur. This can cause a waste of telescope time and can complicate some studies because the accumulated exposure times will be different for different areas in the mosaic. Depending upon the target's ecliptic latitude and the spacing method used, area losses due to overlapping images or under sampling can be up to 30%. For example, one method would be to define the target points with centers at the corners of an equilateral triangle, with sides of dimension of one field of view. Such a pattern will map over 95% of the desired area no matter in what order or orientation the individual fields are made. Judicious dithering will minimize the missing area even further. Efficient mapping strategies will be dependent upon the target's ecliptic latitude and should be developed so to eliminate the need for individual observers to develop them for each project. For example, an observer could specify the desired region and have the system automatically recommend how best to map it and/or present a selection of mapping options and their consequences.

4.9.2 Non-Mapping Studies

For temporal monitoring of a field of objects or the build up of a set of very deep exposures over much of the sky, the primary impact will be the instrument's field rotation causing a loss of the desired area. In a worst-case scenario, a target area near an ecliptic pole over a six week period would incur a 45° rotation in which 17% of the area would be lost if using a square aperture. Thus observers, if needing a fixed area or number of objects for study, would need to scale their requested number of target fields accordingly.

Other DRM programs require the use of a multi-object spectrometer (MOS). As the dispersion direction will be fixed with the instrument and hence, tied to the orientation angle, certain multi-object observations, especially at low ecliptic latitudes, may not be possible. These would be observations where the strongly restrained nominal orientation causes two or more targets to nearly line up in the dispersion direction. Instead of being able to observe these objects simultaneously, they will have to be observed separately, thus potentially increasing the time necessary to make the full suite of observations.

The change in orientation for targets at high ecliptic latitude will also adversely affect multi-object spectrometer observations. If the telescope is allowed to rotate, then there will be a loss of targets (up to the same 17%) between the finder image and spectroscopic observation assuming the same instrument is used to make the finder image and spectra. In addition, unless a position angle for the follow-up observation is chosen (which then sharply restricts when the observation can be made), the follow up observation can occur at virtually any position angle. Thus the observer (or software) will be required to identify which targets are to be observed in the finder image using which position angle. When the follow-up observation is scheduled, then the software must calculate the new orientations of the slits with respect to the instrument field and must determine which targets are still in the field and have positions that will not project spectra on top of other targets.

To maximize the overlapped area and thus the number of targets available for study, would require the followup spectroscopic observations to be made within a few days of the finder observation, or the followup observations must be scheduled in multiples of roughly 90 days (high ecliptic latitudes) or 180 days (low ecliptic latitudes) later and be restricted to the same position angles (–90°, 0°, +90°) and the same small scheduling windows as the mosaic observations.

Attempting to restrict the follow up observations to be within few days would require an observing requirement that forces the finder image to be executed at a maximum off nominal roll and then requires the first visit's data to be downlinked, analyzed, and the follow-up visit updated and executed in as few as 4 days. This would place unwanted (and probably not needed) scheduling restrictions on the second visit. A better plan would be to accept the loss of the corners of the square fields of view of both the cameras and spectrometers for high-latitude fields. Also, there may be some advantage to laying out the focal plan so that the instruments are not at exactly the same orientation, since that configuration will naturally create competition for the same telescope roll angles.

There will be more demand for specific orientations with a multi-object spectrometer than with other designs. Specific orientations will be desired to keep specific target spectra from overlapping, for example, or to keep scattered light from a bright star in the field overlapping the spectrum of a faint target of interest. There will also be observations that desire a specific orientation for a long slit, for example to put a slit along the major axis of a galaxy. For most of the DRM programs, the orientation does not appear to be strict requirement, as long as the orientation is known far enough in advance to configure the apertures on an appropriately chosen statistical sample of objects.

The orientation flexibility will vary with spectrograph design. For spectroscopic observations using slits on particular objects, the micro-shutter and micro-mirror designs may provide an additional ±25° or so of position angle flexibility since a slit and its orientation can be commanded by software. Even more flexible from a scheduling standpoint is the Imaging Fourier Transform Spectrograph, which would allow spectra to be generated at any orientation for any target in its field.

Thus in many cases the observer will be forced to make hard choices: high latitude targets provide position angle flexibility, lower sky background, and longer uninterrupted target availability windows, but have shorter fixed orientation windows, while low latitude targets have restricted position angles, higher sky background (possibly longer exposures), and shorter uninterrupted target availability windows (50 days max. compared to 365 for polar targets), but longer duration at fixed orientations. A fixed orientation constraint would increase the efficiency of certain types of observations but may be impossible to achieve or may cause the observations to be spread out over a time span of months. On the other hand, allowing the telescope to roll may cause an increase in the amount of telescope time required. These options are the extremes and particularly at intermediate ecliptic latitudes the balance between the different effects will need to be made on a case-by-case basis. The operations center will provide tools for observers to determine the effects and weigh their impact on their observations.

4.10 Observation Strategies – Dither Strategies

Most astronomical imaging observations benefit from being made as a series of shorter images, each with a slightly different pointing. This process is called "dithering". The different pointings minimize the effect of hot or cold pixels or other defects on individual chips, as well as the effect of gaps between chips in a given array. It can be reliably assumed that the desire for automatic dithering will be part of NGST's science operations.

Ultimately the issue of whether to dither, and if so by how large an angle and how frequently comes down to a tradeoff in S/N. For the present generation of NIR detectors, detector systematics rather than statistical noise are often the dominant source of error. Sensitivity varies significantly across the face of the detector (and within pixels), and standard calibrations do not always remove all of the detector artifacts. For deep imaging of galaxies the accepted strategy is to observe at a large number of dither positions, spread over an area that is large compared to the largest objects of interest. This allows construction of a sky flat. Typical dithering scales likely for NGST are of order 10-20" on frequencies of once per exposure (e.g. once every 1000 s).

Sub-pixel dithering will be used to improve the resolution of the final image. The requirement is especially strong if the camera is undersampled, which is a design tradeoff being considered in order to maximize the field of view. Even if the camera is Nyquist sampled at 2.2 μm, sub-pixel dithering will be needed to achieve the best spatial resolution at shorter wavelengths. Assuming the detector geometry is well characterised, the sub-pixel offsets do not have to be done within the same physical pixel, and can simply carried out by making the offsets in the larger dither pattern be non-integral multiples of the pixel dimension. In order to do this using the camera as a guider, it must be possible to place the guide star accurately at any location in the pixel.

A common dither pattern (see figure below) moves along a diagonal, with offsets of a non-integral number of pixels. The full dither pattern covers a distance significantly larger than the gap and slightly larger than the largest galaxy. Each of five major steps in the pattern is placed a different sub-pixel offset so that construction of a well-sampled image is possible. The gap is assumed to be 6" and the length of the diagonal dither pattern is ~15". With HST, accurate placement of the target at a given sub-pixel location is not possible for dither patterns this large, but the sub-pixel sampling in those cases comes about because the sub-pixel offsets tend to be randomized. Such a situation may hold for NGST as well, but it would of course be better to be able to choose the exact sub-pixel dither position. Maintaining precisely the same roll angle at the different dither positions is not crucial if the image combination is done on the ground. If it were to be done on board, maintaining the same roll angle (or at least knowing precisely what it is) would crucial because determining relative rolls of two images is computationally expensive.

Figure 5

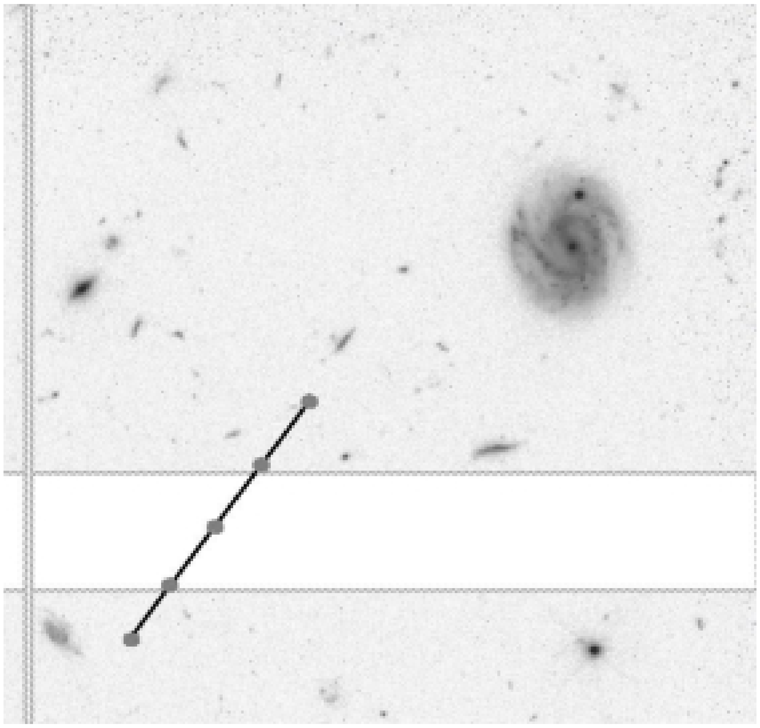

As mentioned in Sect. 4.7, minimizing the effect of the interchip gaps via dithering will greatly increase the chances of the guidestar falling in a gap for at least one of the dithered pointings. Fields with multiple guidestars will be preferred for surveys, and it will be useful to have the ability to switch between them for different dither positions without incurring large overhead.

With the tight constraints on the NGST boresight roll angle, and the projected long exposure times for surveys, the interplay of dithering and survey strategies and orientation becomes quite complicated. It may, for example, be desirable to cycle between filters at each dither position and orientation to ensure that the variation in exposure times across the final image is the same for each band.

It will be useful to have a set of default dithering options for imaging observations, so that the observer need only worry about making a single image of the full instrument's field of view (e.g. 4'x4' for the NIR camera). While it may desirable from the point of view of data transfer to combine dithered images on board, before sending the result to the ground, this is a considerable processing task, and is currently difficult to automate even on the ground. All of the concerns about on-board cosmic-ray rejection and the risks of losing science capability if the algorithms do not work as expected hold for image combination as well.

Spectroscopic observations are probably not free from the desire to dither. While many observing programs could well be tolerant of detector gaps and blemishes, some of the DRM programs are sufficiently long that the orient changes during the multi-day observations will cause the detector gaps to sweep over a significant fraction of the field. In this case it may be desirable to dither to even out the sky coverage. If an aperture or slit is used with the spectrograph, the dithering motion must move the slit as well as the target, against the spectrograph's focal plane. Whether or not this is easy depends on the spectrograph design. The micro-mirror and micro-shutter concepts allow rapid reconfiguration of apertures in software, while larger mechanical devices might imply longer overheads. Even an integral field spectrograph may not be immune from the desire to dither. For example, it may be desirable to place the sources over a range of sub-pixel positions to improve the spatial sampling. Careful attention will have to be given to the stability of the spectroscopic dispersion over the different dither positions.

Mapping observations that require contiguous coverage across the detector gaps will have to live with the reduced exposure times both in the gaps and at the edges of the chips. For the Yardstick camera, the effective exposure time for the final mosaic could be as little as 75% of the exposure time for the best-sampled pixels. There may be an advantage to considering a camera design with gaps almost as large as the chip size. The Cambridge Infrared Survey Instrument is designed this way (Beckett et al. 1996; SPIE vol. 2871, p. 1152). In that case a full mosaic could be constructed from images shifted by the almost the chip size. There may be other advantages to such a design in that the thermal crosstalk between detectors would be reduced. On the other hand, the field of view may not be large enough to support such a configuration.

4.11 Time Scales, Efficiency and Number of Targets

Overheads can strongly drive the type of science program that will be done with NGST. The primary overheads affecting operations will be the slew times, mechanical and thermal slew settling times, momentum management, OTA stability, and possibly readout times. However, currently the overhead details are not well known. For example, the estimates for a 60 degree slew plus mechanical settling range from 20 to 60 minutes. In addition, the thermal settling time after a large change in the sun pitch angle could induce as much as a 2 to 10 hour delay before the image quality is usable.

As a general example, because of target occultations and SAA passages, HST is operated with roughly a 50% efficiency. We would expect to operate NGST at the 75% efficiency level or better. If the efficiency and overhead times are assumed, the number of targets that would be observable in a year is easily calculated. This is shown in Fig. 6. Thus, assuming a 75% efficiency, an average overhead of one hour per observation would imply we could observe about 2200 targets per year and, if desired, execute an HST like science program. However, an overhead of 10 hours per target would imply only 220 targets per year could be observed and would mean restricting the observing program to mainly deep exposures, with each observation a day or two long in duration. Striving for a higher efficiency would place even tighter constraints on the number of targets observable per year.

Envisioning a robust General Observer program for NGST, it is likely that the diverse scientific programs proposed will push toward being able to observe *at least* one target per day without seriously degrading the overall efficiency. Thus the total overhead for long slews and the associated settling must be a small fraction of a day (and consequently that active thermal control of the primary mirror is probably needed).

As discussed in section 4.10, typical integration times will be of order 1000 s, and typical observations will require dithering between exposures. This leads to of order 80 small-angle (< 20") maneuvers per day. The overhead for such maneuvers must therefore be kept down to a minute or two to avoid having the maneuver times drive the observing strategy.

Figure 6

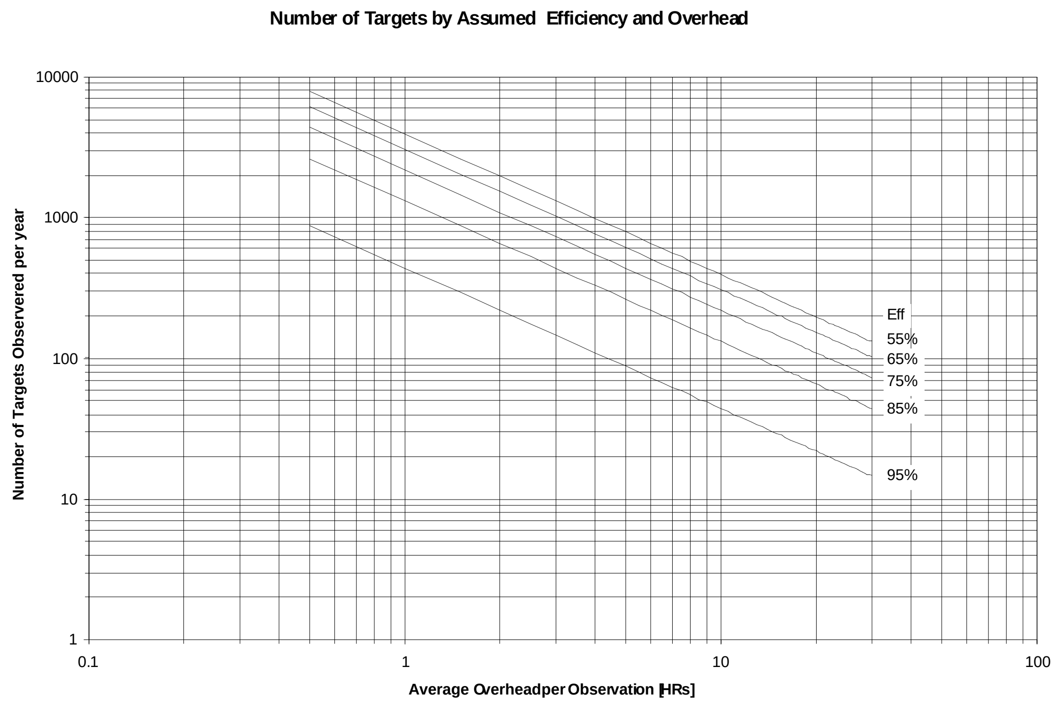

Some of the overheads can and will be mitigated with judicious planning and scheduling of the observations. Minimizing the slew times and the thermal settling times by proper sequencing of targets are obvious examples that will be investigated. If the thermal settling time is the same order as the slew time the telescope could first perform a pitch maneuver and then roll about the sun line to acquire the new target. This will allow the OTA to approach thermal equilibrium during the roll phase of the slew. To minimize the rate of momentum buildup the sun shield needs to remain as close as normal to the sun as possible. This restriction could be implemented by using targets located in a smaller annulus than that given in the yardstick design. Doing this would also minimize the induced thermal changes because there will not be large changes in the solar pitch angle. However, it would also sharply restrict the availability of targets during the year. For example, restricting the pitch motion to keep the sun to within 10 degrees of the sun shield normal would cause most targets to have about 40 days of availability during the year. Another method to minimize the number of momentum dumps would be to schedule back to back observations where the new sun pitch angle would create a torque to eventually cancel the previous momentum buildup. However, this method could increase the OTA thermal instability for a non-symmetric sunshield and would also require the existence of suitable targets, potentially constraining the science program. Studies will be made to determine the best methods once the time scales for the different effects are better known. As the NGST design progresses, it will be useful to develop a detailed overhead budget, and to ensure that no single item comes to dominate the overall budget.

There are changes that will occur to the spacecraft that have time scales on the order of years that can also impact the science observing. It is expected that the insulative properties of the sunshield will degrade with time because of micrometeoroid impacts and other factors. This could cause the rear sunshield temperature to rise by 10 to 30 degrees or more over the life of the mission. This should not affect the NIR observations. However, the increase in the shield temperature during the mission will cause MIR imaging observations scheduled near the end of the mission to take a factor of five or more longer to execute than observations at the start of the mission. This effect combined with the need for a cryocooler for MIR observations implies we should consider a policy to execute a relatively large fraction of the expected MIR observations early in the mission. This policy, when considered, will need to incorporate the relative importance of the MIR observations with the delays that will occur to the NIR science.

All the detectors will degrade with time primarily due to effects of the space environment. However, at L2, NGST will not be protected by the earth's magnetosphere and the rate of damage that will be incurred to the detectors is currently poorly known. The creation of hot or dead pixels, for example, will cause defects on the chips that will need to be removed from the images by more robust (read longer duration) dither patterns.

5 Science Program Options and Observing Models

At this stage in the NGST mission design, it is useful to consider various options for how to define the science program, and ask whether options different from the HST model could result in enhancements to the science or in significantly reduced operations costs. Four different observing “styles” have been considered. All have been used to a certain degree on HST and on other missions. From an operations point of view, the most costly option is to mix all of the different observing styles together. However there may be strong scientific motivations for so doing.

In the following discussion of science and observing models, we first define the various models, and then discuss their costs, benefits, and features. In all cases, the cost largely represents cost to NASA, either through the science center, or through funding of individuals and teams. Features are aspects of a particular model that are either clearly disadvantages (but not necessarily costs) or are not easily categorized as either costs or benefits.

5.1 Science Program Models

General Observer Program (GO)

General observers are selected from the community at large through a widely advertised proposal solicitation, followed by a community-based peer review process. General observing programs are typically allocated a relatively small amount of observing time per program on average, resulting in, e.g. for HST, ~100-500 programs/year. In such a model, individual PIs and their teams are responsible for defining the science goals when submitting the initial proposal (Phase 1 in the HST terminology), for detailed definition of the observations after the observing time is allocated (Phase 2), and for performing their own data reduction and analysis. A typical proprietary time for GO programs (data and intellectual property) is one year.

Guaranteed Time Observer Program (GTO)

Guaranteed observing time is typically provided to the teams who build instruments for a telescope or advise in its development. It is considered an incentive for the long years of effort to develop, test, and construct an instrument, during which time the scientists cannot focus fully on doing their own scientific research. GTOs typically have the privilege of defining their science program and targets before programs are solicited from the general observing community. GTOs are typically given a fixed percentage of the observing time during the early observing cycles, and GTO programs are not externally peer-reviewed. As with the GOs, the GTO team is responsible for detailed program definition in Phase 2, and their own data analysis. The proprietary time is typically identical to that for GO programs.

Key Program (KEY)

Key programs are typically awarded large blocks of observing time (much larger than a typical GO program), and have been solicited from large collaborations, generally multi-institutional. Subjects may be (partially) predetermined (e.g. by a high-level advisory committee), however, science goals could also remain open. Pre-stating a minimum time awarded to key programs encourages large team sizes. If a science program were composed only of key projects, there might typically be 10-50 programs/year. Key projects are peer-reviewed by the same process as the GO programs. The teams are responsible for program development and data analysis, supported by center personnel. Typically a proprietary time would exist. However, for the SIRTF Legacy programs, the pipeline data are to be openly available soon after the observations are taken.

HDF Style Program (HDF)

The HDF style is patterned after the Hubble Deep Field science program successfully executed twice with HST. An HDF style model would involve large blocks of observing time with the science goals defined by community-based panels and other review mechanisms. A science program composed purely on the HDF style model would typically have 10-50 programs/year. The data in the HDF style model would be non-proprietary and funding for research is based on archival/analysis proposals. The unique feature of an HDF style program is that the science center is directly responsible for observation planning, implementation, and data reduction. Calibrated data products would be provided to the general science community within a relatively short time (typically months) after acquisition of the data.

5.1.1 Costs, Benefits, and Features for Science Program Models

The primary cost of a pure component of the science program is related directly to the number (hence size) and fragmentation of the programs, and who (experts or non-experts) is responsible for the resource-intensive efforts of development of the detailed observing program and data reduction. In the table below we highlight the features that are particularly salient as either costs or benefits of a particular style of science program. The table also identifies more cost-neutral features that distinguish one style from the others either in terms of scientific productivity, or in terms of operational implications.

Model	
GO	***Benefits:*** Engages entire community Encourages diversity of ideas, creative thinking in a large community Highly competitive; Gives TAC many options from which to choose Allows short programs More responsive to changes in scientific questions than a series of long-term projects ***Costs:*** Large number of proposals; expensive TAC process Large number of non-expert users; increases requirements for user-friendly software, documentation, user-support personnel, etc. Diversity of observing strategies, science goals leads to more complex data-reduction software. Need to balance competing requirements among a large number of programs and teams Adds some complexity to the scheduling process. ***Other Features:***

	Lacks coordinated community effort to arrive at consensus on key scientific questions. Large oversubscription tends to make proposers shy away from ambitious proposals. Small, short-term grants imply more administrative overhead. More pressure to propose for funding reasons rather than science reasons. Generates diverse, heterogeneous archive
GTO	***Benefits:*** GTO time provides incentive for instrument team to develop the best instrument. Less documentation and user-support required. A Time Allocation Committee is not used, resulting in cost savings . Smaller number of observers than GO model reduces user support costs. GTO teams contribute to on-orbit calibration and characterization of their instruments. Fewer grants means less administrative overhead than GO model ***Costs:*** Procedures need to be developed to protect GTO programs from duplication by GOs Tendency among GTOs to "push the envelope" of instrument capabilities increases costs associated with command development and calibration of specialized instrument modes. Science-team costs drive up the phase C/D costs of the instruments. Science team costs are significantly higher than GO costs ***Features:*** Research projects proposed at the time of instrument selection may be irrelevant by the time of launch. Connecting the GTO time award to performance in meeting schedule, budget, specs might be better than the HST model. GTO archive is likely to be less diverse than a GO archive
Key	***Benefits:*** Major scientific questions are addressed with ambitious programs. Key projects help highlight mission goals to a broader community. Smaller number of observing programs reduces user support costs. Fewer proposals reduces proposal solicitation, TAC costs. Fewer grants means less administrative overhead than GO model. ***Costs:*** Need for user-friendly software, good documentation nearly as high as for GO model. Telescope efficiency may be lower if the overall observing program lacks short programs to fill in gaps. Science team costs typically somewhat higher than GO costs for an equivalent amount of observing time. ***Features:*** Less overall community involvement in NGST. Difficult to allocate time to short projects. Both successes and failures of the projects are highly visible Multi-year projects risk becoming irrelevant as time progresses. Allocating a large-fraction of the telescope time to long-term projects makes it harder to respond to new scientific opportunities. Archive is highly focused, but deep
HDF	***Benefits:*** Encourages multi-purpose observing programs. Encourages surveys to support key science projects. Involvement of operations center in planning, execution, and reduction builds expertise that can be applied to other observing models. A pure HDF model would reduce costs of developing user-friendly software and

documentation.
Direct user-support costs minimized.
Observing efficiency and scientific goals can be traded off more explicitly.
Cost of external committees comparable to Key Project TAC costs.

Costs:

Telescope efficiency may be lower if the overall observing program lacks short programs to fill in gaps.

Features:

Risk that projects will lack clear scientific leadership.
Less overall community involvement.
Difficult to allocate time to short projects.
Both successes and failures of the projects are highly visible.
Multi-year projects risk becoming irrelevant as time progresses.
Allocating a large-fraction of the telescope time to long-term projects makes it harder to respond to new scientific opportunities.

If only one of the above science program models were to be adopted, the pure HDF-style model is likely to be the least expensive. The major savings are due partly to reduced requirements on software and documentation, and partly to the elimination of pre-observation support for general observers. The other three models would all require more software and documentation. The GO model, with the largest number of observers to support, is likely to be somewhat more expensive than the Key-project or pure-GTO models.

It is of course likely that several, if not all, of the above observing models will be adopted for NGST. All are familiar from other national and international facilities, and all have their own scientific benefits. Nevertheless, employing such a mixture of science observing modes is likely to be more costly than sticking with one science-program model. If operations costs become a serious driver, it may be that the most savings could be realized by delaying the start of the GO program, so that the costs of developing software and documentation can be spread out. Calibration, software, and documentation will be more mature by the time the diverse GO community begins to use the telescope. On the other hand, early involvement in NGST from a broad community is sure to be scientifically worthwhile, so the delayed-GO option (modeled after SIRTF) is not necessarily the best overall.

5.2 Observing Models

Without constant interruptions from earth occultations and SAA passages, the NGST observing schedule should be much simpler than HST's. An inherent complication for HST scheduling is the fact that the orbit evolves significantly over a period of months. Precise start and stop times cannot be specified by the observers far in advance because the target availability is not known precisely enough at the time of Phase-2 proposal submission. Observers typically specify constraints on the target observing window, either explicitly by imposing timing requirements or implicitly by imposing roll-angle, low-sky, continuous-viewing-zone, or other such requirements. It is then left to STScI planners to juggle the constraints of different observing programs to try to maximize the overall observatory efficiency.

For NGST, target availability will generally be known a long time in advance, making it possible in principle to give observers a fixed time on the observing schedule at the time of proposal selection. From an observer's point of view, removing the uncertainty of the scheduling would remove one of the great frustrations of using HST. However, the instability of the HST observing schedule has come primarily as a result of instrument problems, changes of servicing mission launch dates, the appearance of targets of opportunity, and other events that caused rescheduling. NGST will probably not be immune from such interruptions. Servicing missions will not be an issue, but solar flares will cause major interruptions on relatively short notice. The NGST L2 orbit will not be a panacea for all the ills of the HST schedule.

There appear to be two major policy choices to be made. The first is whether to schedule HST observations in a queue, or in fixed calendar blocks. At some level the "block" concept will exist in any operations scheme; the main question is what unit to use for the block size (e.g. hours, days, or weeks – for HST it is orbits), and whether the schedule will involve wholesale interleaving of blocks from different proposals as for HST. The second policy choice is how to respond to interruptions or disruptions of the schedule (i.e. whether to reschedule the observations that were displaced). For HST, observations lost due to circumstances beyond the observer's control are typically rescheduled. For ground-based observatories, they are typically lost, and the observer must re-apply. There may be significant cost savings in adopting the ground-based approach.

Four possible observing models have been examined for their applicability to NGST:

- Block scheduling/real-time ops the norm (e.g. ground-based observing, IUE)
- Block scheduling/real-time ops the exception (e.g. VLA)
- Queue scheduling/real-time ops the exception (e.g. HST)
- Mission style observing teams (science teams do the scheduling, e.g. Astro-1,2, planetary missions)

For most practical purposes, "mission style" observing is simply block scheduling in large blocks, with large interdisciplinary science teams. The unique aspect of this model is that overall scientific optimization is a greater consideration in the constructing the schedule. Constraints and even science goals of different programs are explicitly traded against the practical aspects of scheduling the observations. However, while the scientific interactions in scheduling are different, the cost of mission-style scheduling will be comparable to block scheduling.

Similarly, the costs and benefits of real-time operations (where the observer is present at the control center and interacts with the telescope operators) are essentially independent of whether it is part of block or queue scheduling. The primary benefit is the ability to review the data and modify the observing strategy part way through the observations. The primary costs are in telescope operators, robust real-time commanding software, and the necessity for nearly continuous communications. If real-time operations were the norm, there might conceivably be some cost savings associated with reduction in the complexity and documentation of phase-2 proposal preparation software, but this savings will be largely offset by the cost of real-time commanding software and personnel, and may come with a scientific cost of having less optimal proposals at the time of execution.

Given that real-time operations and "mission style" observing appear to be either cost-neutral, or possibly cost- and value-added, options, we focus the discussion that follows on the pros and cons of block vs. queue scheduling. We have tried to define the two options in a way that makes them clearly distinct, but nevertheless both capable of scheduling the likely variety of NGST science.

5.2.1 Block Scheduling

In the block scheduling model, observers (except those requesting targets of opportunity) are allocated observing time on fixed calendar dates. The assignment of dates could be done as part of the TAC process, or could be done by the long-range planning system shortly after TAC allocation. This model would be difficult to implement with HST for several reasons:

- Typical observations are only a few hours.

- Many observing programs involve scheduling large numbers of targets for short observing times.
- Due to evolution of HST orbit, target availability windows and times of SAA passages are known only approximately when the time is allocated.

Block scheduling is familiar from ground-based observing. The tradition at ground-based observatories has been to leave the detailed scheduling of exposures, and to a certain extent even targets, largely to the discretion of the observer, who is present at the telescope when the observations are carried out. IUE was also scheduled in blocks, with the observer present during the observations to allow real-time decisions to be made. The VLA model falls somewhere in between block and queue scheduling. While the schedule is put together well in advance of the observations, it is not part of the TAC process, and observations from different proposals are interspersed to take advantage of target availability. The proposer is often not present when the observations are taken.

A few ground-based observatories are experimenting with queue scheduling. This is driven largely by the desire to optimize the use of diverse observing conditions. Queue scheduling allows observations that need good seeing to execute when the seeing is good, for example, and for other programs with less stringent requirements to fill in when the seeing is poor. Periods of good seeing are not predictable. As is the case for low-earth-orbiting satellites, queue scheduling is an attempt to allow scientific optimization in the face of observing conditions that cannot be predicted at the time of selection.

Apart from unpredictable instrument safings and shut-downs due to solar flares, NGST observing conditions can be predicted far in advance. It may thus be practical to consider allocating fixed calendar blocks of observing time. There are several potential advantages to this model:

- Observers know the time of observation far in advance and can plan accordingly. This is particularly important for NGST spectroscopic observations, where target selection will depend on the roll angle. It will also be of some importance to NIR imaging observations requiring low zodiacal background, or MIR observations requiring low thermal background. Requirements for specific orientations and background would have to be specified as part of the initial proposal.
- Scheduling optimization could (in principle) be part of the TAC process. Since target availability can be determined prior to proposal selection, part of the time-allocation process could be to ensure that there are feasible observations for NGST for every day of the year, or at least ensure that large programs do not conflict in their scheduling requirements.
- Real-time observations, if any were to occur, would be more practical than in a queue scheduling model.
- Observers would have more control over the optimization of their observing time and the sequencing of their observations. They would also in principle know the timing of their observations relative to calibration observations, and could respond by inserting their own calibrations if they felt it necessary.
- With block scheduling there is no need to oversubscribe a cycle of observations. However, unless all of the observing constraints are known by the TAC and the scheduling conflicts resolved as part of the time-allocation process (which is probably impractical), the price to pay for this certainty of scheduling is that the observations may not be scheduled at the scientifically optimal time. The analogy from ground-based observing is getting allocated time when the target you proposed for is only available half the night --- as frequently happens.

5.2.2 Queue Scheduling

Queue scheduling is familiar from HST and other low-earth-orbiting satellites. Observations are not assigned specific times when proposals are approved. The total amount of time allocated by the TAC is based on an estimate of the average telescope efficiency, and the TAC does not consider whether that efficiency can actually be achieved for the set of targets and observations they have approved. Observations are placed in a pool, and the schedule is set up to optimize the overall observatory efficiency. Observations within a given proposal are re-ordered and interspersed with observations from other proposals as part of this global optimization. For HST, the specific timing and orientation requirements are not specified in the phase-1 proposal considered by the TAC, but are specified later as part of the phase-2 process for approved proposals. There is frequently some iteration between the observer and program coordinator on the exact specifications of the observation. It is often a challenge, for example, to find a scientifically desirable orientation that has viable guide stars and can schedule when the target is observable.

Almost by necessity, queue scheduling implies some amount of oversubscription. If the schedule is not put together until well after the selection process, and if observers are not forced to observe at times that may not be optimal for their programs (as would

There are several potential advantages to queue scheduling for NGST:

- Exact specification of the observing requirements does not have to be done at the time of proposing. This saves proposers a lot of effort (which would be largely wasted effort, since most proposals are rejected).
- The TAC would not have to consider scheduling issues.
- NGST slew times could be minimized.
- Unpredictable housekeeping activities could be inserted into the queue with less noticeable impact.
- Targets of opportunity could be scheduled without explicitly displacing another observer's block.
- Mission planning is more robust against observer error or responsiveness.

The choice of block vs. queue scheduling in many ways hinges on what it means to have a scientifically efficient observatory. If observing constraints are not considered at all up-front in accepting proposals, then to make an efficient schedule one must either oversubscribe (accept more proposals than can be scheduled during the cycle) or force observations to occur at times that are not entirely optimal (e.g. have a non-optimal orientation). Solving the scheduling problem by oversubscription means that proposals of lower scientific priority may be scheduled in favor of proposals of higher scientific priority. Solving the scheduling problem by block scheduling implies that observers may have to make scientific compromises in order to accommodate non-optimal observing conditions. Block scheduling does not require oversubscription but implies scientific compromises at the proposal level. Queue scheduling either requires oversubscription, in which case some accepted proposals will not be executed, or requires that observers accept non-optimal observing conditions that they find out about at a much later date. One form of oversubscription, which may alleviate some of the difficulties with queue scheduling, is to have a "filler" program akin to HST snapshots. There are varied opinions on whether or not this is a scientifically efficient use of telescope time.

5.2.3 Policies for Missed or Displaced Observations

The choice of which scheduling model is best for NGST may depend to a large extent on what the policy is for observations that are missed or degraded due to instrument problems or solar flares, or observations that are displaced in favor of targets of opportunity. The rate of observation failures on HST is typically low (a few percent), but the replanning rate is much higher due to changes in servicing mission dates, or changes in science instrument behavior. The situation is likely to be slightly different for NGST. Assuming the instruments are relatively stable, the major interruptions will be due to solar flares and coronal mass ejections, which could impact up to ~20% of the observations during solar maximum.

In the HST model, observations are automatically rescheduled (nearly always), and do not by default go to the end of a queue, but rather back into a pool. The HST schedule is thus frequently updated to respond to changing conditions of the spacecraft or instruments and a changing pool of programs remaining to be executed. There are both advantages and disadvantages to this approach. The main advantage is that observers get their data, eventually. The primary disadvantage is that they cannot predict exactly when they will get their data, and sometimes it is so long after the observations were proposed that the observations have lost some of their scientific relevance.

The issue of rescheduling observations will exist for NGST, but may not be as large a problem as it is for HST. The HST servicing missions, for example, are a major source of schedule disruptions, and these will not exist for NGST. For NGST, the major disruptions will be instrument safings and solar flares. Targets of opportunity will likely result in only a few displacements. There are multiple options for how to respond to missed observations:

Policy	Advantages and Disadvantages
Model 1. Treat disruptions as "weather." Observers are required to re-apply for time.	***Advantages*** Familiar to ground-based observers. Minimizes schedule disruption. Minimizes time between proposal and execution for observations that were not disrupted. Delayed proposals can be reconsidered by the TAC in the light of new scientific developments.
	Disadvantages Observers forced to re-propose. Protests are likely when the disruption is not clear-cut (e.g. poor instrument performance). Rigid application of the policy would undoubtedly mean that some high priority observations (almost certain to get time in the next TAC) would be delayed while lower-priority observations remain on schedule. Exceptions to the policy may be needed for • Observations displaced by a target of opportunity • Displaced calibration observations • Displaced observations that are part of a linked sequence.
Model 2. Reschedule observations to the end of the queue	***Advantages*** Minimizes schedule disruption. Observers not forced to re-propose.
	Disadvantages Maximizes time between proposal acceptance and execution for disrupted observations. Scientific priority does not enter into the rescheduling decision. Targets displaced from one cycle affect the scheduling of the next

	cycle. Thus the long-range plan for the next cycle cannot be stable until the previous cycle finishes. Exceptions to the policy may be needed for • Displaced calibration observations • Displaced observations that are part of a linked sequence.
Model 3. Reschedule observations as soon as possible in a tightly scheduled queue	***Advantages*** Can keep disrupted or lost observations close to the original schedule. Observers not forced to re-propose. Scientific priority can play a role in deciding where to reschedule the observations, and which observations they may displace. Timing and sequencing constraints can still be met for the disrupted observations.
	Disadvantages Maximizes schedule disruption for observations that would not be otherwise impacted by the schedule interruption. If there are a lot of interruptions, the observing plan will need to be updated frequently. Timing and sequencing constraints for the disrupted observations will need to be weighed against timing and scheduling constraints for the observations they will displace. Programs without tight timing or sequencing constraints will tend to get bumped to the end of the queue.
Model 4. Reschedule observations as soon as possible in a queue that includes some low-priority "filler" or "supplemental" observations, or long-term surveys that can be stretched out.	***Advantages*** Preserves the advantages of rescheduling, but with less disruption to high-priority programs.
	Disadvantages Default observing plan contains low-priority observations. Even with filler observations, it may be necessary to shuffle around high-priority observations to re-insert disrupted observations with tight timing constraints.

5.2.4 Costs

Both queue and block scheduling require software tools to predict target observation windows, guide star availability, observatory overheads, etc. Both require expert support to assist observers in preparing their observations. Both require software for optimizing slew time. The major difference between the two models is when and how the global optimization is done. Queue scheduling allows the optimization to be done later, and even multiple times, and allows sections of different proposals to be interspersed to increase observatory efficiency. Cost may not be a major factor in deciding which model to adopt, or whether to adopt some variant in between.

There may be more significant cost differences associated with the policy for how to reschedule missed or disrupted observations. The least expensive mode of operation is to force the observers to re-apply for time. For a rough estimate of *relative* costs, assume the following:

- Reinserting an observation into an existing queue (and rescheduling the observations it displaces) requires one "work unit" (WU; likely to be somewhere between one FTE day or one FTE week, depending on the complexity of the schedule and observing constraints).
- Costs of moving an observation to the end of the queue, or deleting it entirely are insignificant.
- A scientific or technical iteration on a proposal due to rescheduling, even at the end of the queue requires one WU (an example would be changing a slit mask for a different orientation).
- 1000 target visits per year (roughly 10% of which are calibration).
- 15% of the visits in a given cycle are lost or disrupted for one reason or other. (This is roughly the rate expected if the observatory is shut down during periods of high particle background during solar maximum.)
- 20% of the rescheduled visits would require a scientific or technical iteration.

Under these assumptions the cost of rescheduling in each of the above models is:

Model 1	18 WU	Assumes 15 calibration observations will need to be rescheduled, and 3 of those will require scientific iteration.
Model 2	48 WU	Assumes the above 18 WU for rescheduling calibrations, plus 150 x 0.2 WU for the scientific iterations required on 20% of the 150 observations bumped to the end of the queue
Model 3	108 WU	Assumes each rescheduling displaces on average ONE other observation already in the queue, and that 20% of the total (rescheduled + displaced) would need scientific iterations. 18 WU for calibration rescheduling + 30 WU for rescheduling 150 observations 60 WU for scientific iterations on the proposals (rescheduled + displaced)
Model 4	81 WU	Assumes each rescheduling displaces on average 0.1 other (non-filler) observations already in the queue. 18 WU for calibration rescheduling 30 WU for rescheduling 150 observations 33 WU for scientific iterations

The cost of model 1 is roughly 20% of the cost of model 3. The costs of rescheduling are not the major cost in scheduling the observations, so cost may still not be the decisive factor in deciding what policy to adopt.

5.3 Proposal Cycle Models

The primary goal of a proposal solicitation process is to select the best possible science program, ensuring results of unquestionable importance which change our understanding of the Universe. Secondary goals include: "buy-in" from the community that the selected program is the best and that the selection process is open and fair; optimal cost to benefit, meaning that the selection process should not require excessive resources either from STScI or the community; and if on-site reviews are held, valuable visibility from the community into STScI's role and expertise, and valuable feed-back from the community to us.

The typical proposal cycle model for most space-based observatories with general observing programs has involved an open peer-review process wherein any program of any length on any topic is solicited on a regular basis, typically one year. This model has a number of disadvantages that motivate consideration of other models. First, the length of time between idea and execution is often long (one year at best, but as long as two years, and sometimes even longer, if unexpected problems crop up). Second, the number of proposals that need to be reviewed at one time is large, requiring a large expenditure of resources for a short time on a cyclical basis. Third, assembling and managing a large proposal review is costly and time-consuming. Fourth, while an open, HST-like proposal review would presumably attract many proposals, a large oversubscription can alienate the losers in the proposal sweepstakes. Funding is also a particularly important issue. In the traditional approach relatively few people get large amounts of money.

One goal in considering other models is to shorten the time between idea and execution. This can be achieved in either of two ways: by using a shorter, fixed cycle length (e.g., quarterly or semi-annually); or by using a "rolling" proposal cycle. Examples of both types of models already exist within the astronomical community. Ground-based observatory proposal solicitation commonly occurs on a semester or quarterly basis, and NSF proposals are accepted on a rolling basis without a fixed, annual proposal deadline.

While a quarterly review would indeed shorten the time from idea to execution, given the proposal pressure that exists for large, space-based telescopes such as HST, it is not clear the proposals would distribute themselves smoothly in a quarterly fashion, or be submitted in the appropriate quarter based on their schedulability. On the other hand, if observers were aware in advance of the precise windows within which their observations could be scheduled, thus naturally dictating the quarter or half of the year to submit proposals, it might smooth out the annual workload variations.

In a rolling review type model, proposals would be accepted and reviewed continuously rather than in response to a deadline. The ideal would be to have a core set of programs (e.g., the GTO, key projects, and a "pool" of smaller programs) and achieve steady state with a pool of observations always available for execution. The advantages of this approach include a much shorter interval between science idea and execution, and thus a fresher science program. It also spreads out the workload in a natural way, eliminating the seasonal cycles of activity. It eliminates the difficulty of planning in advance of unplanned outages, since it is likely there will always a small pool of executable programs, continuously replenished. It eliminates carryover, since there is never a large pool.

The rolling model has several disadvantages, including the increased difficulty of achieving a science program that is well balanced among different types of observing projects and teams, and ensuring a consistent review of all proposals. The latter could be addressed by maintaining a standing panel over the course of a year that ensures consistency and understands the program balance. However, it is important to ensure that the process is fair, and is perceived to be fair by the astronomical community. With a rolling proposal cycle it may be difficult to ensure that enough proposals are always available to allow efficient scheduling, especially if NGST has limited scheduling flexibility.

6 Backend Systems

The NGST backend data processing, archiving, and data distribution systems will be conceptually similar to the backend HST systems. Standard processing – data reformatting and calibration – needs to be done uniformly for all data in order that the body of NGST data has maximum utility for archival research. It is also more cost effective to have a common approach to calibration rather than to leave development of calibration algorithms and codes completely in the hands of the community. The latter approach leads to massive duplication of efforts and inconsistency of results.

6.1 Pipeline Processing

Data from the NGST will be downlinked in a telemetry format, most likely including some form of data compression in order to optimize use of the downlink bandwidth. The pipeline processing system will support functions similar to the HST pipeline system:

- Telemetry decommutation
- Data decompression (as necessary)
- Population of keywords from associated engineering information and proposal information
- Linking and co-processing of associated images (CR splits, dither patterns)
- Standard calibration processing
- Population of keywords with results of calibration (e.g., statistical measures)
- Data quality assessment – did the data successfully complete all processing steps in the pipeline?

With NGST we have an opportunity to provide a much higher level of data product – object catalogs. The majority of the programs in the DRM and the expectation for much of the GO science will be based on long exposure times and standard CR-split and dither patterns. With additional moderate constraints or requirements on observing programs one could assure that all NGST data is amenable to higher level processing and automated source extraction and cataloging. The resulting source catalogs would be an important complement to the large all-sky surveys that are being carried out in the early 2000s: 2MASS, POSS-II/GSC-II, and SDSS. Thus, to the list of HST-like pipeline processing functions, for NGST we add, as a desirable goal

- Source extraction and object catalog construction

Depending on the stability of NGST calibrations it may be necessary to do source extraction and catalog construction downstream from the initial pipeline processing. Nevertheless, the catalog construction process itself would be done in the context of a highly automated pipeline. The cataloging process could also include automated spectral extractions, analysis, and classification.

6.2 Archiving

The NGST data archive will be similar in structure to the HST and other MAST archives, with the important addition of an object catalog. Users will be able to search for data as they do now for HST – by target name, target coordinates, observation time, PI, program ID, or any combination of various instrumental parameters. They will also be able to browse and mine the object catalog and easily cross-correlate it with other catalogs and archives, both within MAST and at other astronomy data centers. By ~2010, when the NGST archive begins to be substantially populated, we expect that NASA and NSF observatories will be participating in a confederation of data services that support full interoperability. That is, users will be able to easily query across data centers for information related to an object or class of objects, and to effectively mine catalogs of tens of millions of objects to find those that can be cross-identified or that occupy specific or unusual areas of parameter space. The NGST archive must be a full participant in this confederation.

The baseline data rate from NGST is similar in magnitude to the HST data rate with ACS in operation (e.g., an average of 80 1000-second exposures with an 8k x 8k imaging detector

= 10 GB/day). However, current studies of the cosmic ray impact rates indicate that 1000-sec exposures may be too long, resulting in unacceptably high levels of data loss due to CR events. During solar maximum even shorter exposures (less than 300-sec) may be required to keep CR contamination to levels below 10%. Total daily data rates may be in the range of 10-100 GB, corresponding to annual archival data volumes of 2-20 TB. Data storage systems are now available at moderate cost for on-line archives up to ~10 TB in size. In the coming decade advances in mass storage technologies will enable on-line archives of the scale of 100s of TB. Our strategy for NGST will be to select an appropriate storage technology at the latest possible time (in order to maximize storage and minimize cost), and be prepared to migrate to new technology storage at least once in the NGST mission lifetime.

6.3 Data Distribution

With a standard image size of 128 MB (non-destructive read-outs combined), NGST data, if available now, would not be conducive to network transfers without use of lossy data compression. Lossy compression algorithms, however, may be appropriate for NGST images in which the majority of image pixels are blank sky. Lossy compression algorithms now exist which preserve full fidelity in image regions containing significant signal, yet achieve overall compression ratios of 10:1 to as much as 100:1. We also expect wide-area network bandwidths to increase substantially in the coming decade, to the point where downloads of many GB of data may be feasible in acceptable times.

Hard media distribution will also be an option. By the time NGST is in orbit we can expect a successor technology to DVD (which currently supports 3 GB of data per disk) with a capacity of perhaps 20-40 GB per unit. Such media could support the distribution of 100-200 NGST images at full resolution. Lossless compression would increase this number by a factor of 2-3, and lossy compression by a factor of 10 or more.

The real benefits of NGST in terms of cross-correlative archive use will require real-time network access. Archival researchers are likely to focus on the high level data products such as the object catalog, rather than going back to original data (much as was observed with the community use of the Hubble Deep Field observations).

7 Summary and Issues for Future Study

This document has presented an initial overview of NGST operations issues and tradeoffs. A major goal has been to identify the design decisions that might have a strong effect on operations, and the operations issues that may have a major bearing on the NGST design. In priority order, from an operations perspective the key issues are:

1. Data Volume
2. OTA stability and calibration
3. Thermal crosstalk in the ISIM
4. Overheads for small slews & dithering
5. Telescope boresight roll restrictions
6. Guide camera
7. Thermal stability and time to settle after slews
8. Momentum management
9. Radiation environment

Not all of the technical issues have been studied by the operations group, and for at least some the Yardstick design may no longer be the most useful benchmark to consider. There are, however, a few clear recommendations that can be proposed at this time.

Recommendations:

1. Data Volume

   NGST should have a downlink capacity of 250 Gbits/day and an on-board storage capacity of 135 Gbits. Because this is expensive, further study of on-board cosmic-ray rejection with realistic assumptions for the detectors is needed.
2. OTA stability and calibration

   The OTA stability and how it is to be maintained could have a major impact on operations. For example, a requirement to minimize changes in the sun pitch angle to maximize thermal stability would impose a major constraint on the scheduling of observations. Also frequent PSF monitoring and OTA adjustment could be major overheads. Further study is needed. Specific schemes of OTA management need to be worked through to the point where operational impacts are understood.
3. Thermal crosstalk in the ISIM

   Operations are greatly simplified if the operation of one instrument or subsystem does not affect another. Meeting the NGST requirements on thermal stability poses a serious technical challenge. If one of the technical solutions were to prevent certain kinds of parallel operations (e.g. readouts of one detector while observing with another), it would impose additional scheduling constraints that would have to be met by the ground-scheduling system and the Observation Plan Executive. A systems engineering study is needed. If scheduling constraints are going to be part of a solution, operations issues, and the viability of event-driven operations, would need detailed study.
4. Overheads

   The combined total of overheads, including those for slews and small-angle maneuvers, momentum management, wavefront calibration, detector readout, and instrument reconfigurations, must be kept low, and any individual component should be kept small enough that it does not have to be a key parameter in the scheduling algorithm. An overhead budget should be developed and maintained as the NGST design progresses to ensure that scheduling efficiency and complexity can become part of the cost-benefit tradeoff
5. Roll Restrictions

   Increasing NGST boresight roll flexibility could greatly simplify scheduling. If the off-nominal roll capability for most targets is less than ~10 degrees on any given calendar day, the choice of field orientation for a specific observation will become an operations choice rather than a scientific decision.
6. Guide camera

   Apart from the technical issue of whether the guide camera can really provide the sampling rate needed for guiding, there are major operational and scientific issues to consider. Survey strategies will depend on the availability of guide stars, both for the prime position and for the various dither positions. If 1/4 of the pixels are dedicated to guiding that will affect S/N calculations. As the boresight roll angle changes the locations of the three remaining camera fields will rotate around the guide star (if roll restrictions are tight), further decreasing the contiguous area that can be observed for long exposure times or with multiple filters. Finally, the ability of the guider to place a star accurately at a desired sub-pixel location will affect some observing and acquisition strategies. Further study of the guide-camera option and the impacts of other guide designs are needed.
7. Radiation environment

   The design and complexity of the scheduling system will depend in part on the philosophy for rescheduling observations that are missed or degraded due to particle radiation or other problems. Early policy decisions could help guide the design of the scheduling system.

8 References

"Controlling Radiation Torques for 3-Axis Control", NGST Technical Report, 23 June 1997, by Lloyd Purves and Peter Stockman

"Analysis of Solar Pressure and C.M.-C.P. Offset on Momentum Unloading", Memorandum, 17 July 1997, by Gary Mosier

"Space Mission Analysis and Design", 1991, edited by James R. Wertz and Wiley J. Larson.

"Solar Proton Flux -- Analysis of GEOS SEM Data, Nov. 17, 1999" Isaacs, J.,. http://presto.stsci.edu/vision/vision_se/ngstOpsConceptWG/CommunicationsStudy/index.html

"Draft Communications Study Report, October 13, 1999," Isaacs, J., http://presto.stsci.edu/vision/vision_se/ngstOpsConceptWG/CommunicationsStudy/index.html

"Analysis of Solar Pressure and C.M.-C.P. Offset on Momentum Unloading", Memorandum, 17 July 1997, by Gary Mosier

"Controlling Radiation Torques for 3-Axis Control", NGST Technical Report, 23 June 1997, by Lloyd Purves and Peter Stockman

Spagna 1999; http://ngst.gsfc.nasa.gov/public/unconfigured/doc_458_2/IRCounts.pdf

The Next Generation Space Telescope: Visiting a Time When Galaxies Were Young, Stockman, H. S., ed. 1997, STScI

"Space Mission Analysis and Design", 1991, edited by James R. Wertz and Wiley J. Larson.

Ops Concept "Primer sessions," http://presto.stsci.edu/vision/vision_se/ngstOpsConceptWG/primer/index.html

The NGST Design Reference Mission

9 Acronym List

2MASS	2 Micron All Sky Survey
ACS	Advanced Camera for Surveys
C&DH	Command and Data Handling
CPU	Central Processing Unit
CR	Cosmic Ray
DRM	Design Reference Mission
DVD	Digital Video Disc
FGS	Fine Guidance Sensor
FOV	Field of View
FSM	Fast Steering Mirror
FTE	Full Time Equivalent
GB	GigaByte
GEOS	Geosynchronous Operational Environmental Satellites
GO	Guest Observer
GSC-II	Guide Star Catalog II
GSFC	Goddard Space Flight Center
GTO	Guaranteed Time Observer
Gbits	GigaBits
HDF	Hubble Deep Field
HST	Hubble Space Telescope
Hz	Hertz
I&T	Integration and Testing
ID	Identification
IR	Infrared
ISIM	Integrated Science Instrument Module
IUE	International Ultraviolet Explorer
InSb	Indium Antimonide
L2	Second Lagrange Point
MAP	Microwave Anisotropy Probe
MAST	Multimission Archive at STScI
MB	MegaByte
MIR	Mid-Infrared
N-m	Newton-meter
N-m-s	Newton-meter-second
NASA	National Aeronautics and Space Administration
NGST	Next Generation Space Telescope
NIR	Near-Infrared
NSF	National Science Foundation
OPE	Observation Plan Executive
OTA	Optical Telescope Assembly
PC	Personal Computer
PI	Principle Investigator
POSS-II	Palomar Observatory Sky Survey II
PSF	Point Spread Function
QE	Quantum Efficiency
REE	Remote Exploration and Experimentation
S/N	Signal to Noise
SAA	South Atlantic Anomaly

SDSS	Sloan Digital Sky Survey
SEM	Space Environment Monitor
SI	Science Instrument
SIRTF	Space Infrared Telescope Facility
SSM	Spacecraft Support Module
STScI	Space Telescope Science Institute
TAC	Time Allocation Committee
TB	TeraBytes
TDRS	Telemetry and Data Relay Satellite
URL	Uniform Resource Locator
VLA	Very Large Array
WFPC2	Wide Field Planetary Camera 2
WU	Work Unit